arXiv:1702.02227v1 [math.NA] 7 Feb 2017

Inverse regression for ridge recovery

Andrew T. Glaws[a], Paul G. Constantine[a], R. Dennis Cook[b]

[a]*Department of Applied Mathematics and Statistics, Colorado School of Mines, Golden, CO 80211*

[b]*Department of Applied Statistics, University of Minnesota, St. Paul, MN 55108*

Abstract

We investigate the application of sufficient dimension reduction (SDR) to a noiseless data set derived from a deterministic function of several variables. In this context, SDR provides a framework for *ridge recovery*. A ridge function is a function of a few linear combinations of the variables—i.e., a composition of a low-dimensional linear transformation and a nonlinear function. The goal of ridge recovery is: using only point evaluations of the function, estimate the subspace that is the span of the ridge vectors that define the linear combinations. SDR provides the foundation for algorithms that search for this *ridge subspace*. We study two *inverse regression* methods for SDR—sliced inverse regression (SIR) and sliced average variance estimation (SAVE)—that approximate the ridge subspace using point evaluations of the function at samples drawn from a given probability density function. We provide convergence results for these algorithms for solving the ridge recovery problem, and we demonstrate the methods on simple numerical test problems.

Keywords: sufficient dimension reduction, inverse regression, ridge recovery, ridge functions

1. Introduction

Advances in computing have enabled complex computer models of physical systems that complement traditional theory and experiments. An experimen-

Email address: aglaws@mines.edu (Andrew T. Glaws)

talist supplements physical measurements with statistics that quantify random noise and bias in the experimental system. The need for analogous quantification in computer model predictions has led to the field of *uncertainty quantification* (UQ) [1, 2, 3]. In contrast to a noisy experiment, the computer model is a deterministic map from physical inputs to outputs, where the outputs provide the model's predictions. Forward UQ treats the inputs as random variables and estimates the resulting outputs' probability measure or associated statistics; inverse UQ estimates a probability measure on the inputs such that the resulting output measure is consistent with given experimental measurements. Accurate estimation of statistics requires many computer model runs, which may not be feasible when a single model run is computationally expensive. In this case, one may construct a cheap surrogate model—e.g., a polynomial or radial basis function approximation—to use in place of the expensive computer model. However, such constructions suffer from the *curse of dimensionality*—i.e., the cost constructing an accurate approximation increases exponentially as the number of input variables increases [4]. Therefore, quantifying uncertainty in expensive computer models with several input parameters remains a difficult challenge.

One approach to address this curse is to seek a low-dimensional parameterization of the input/output map such that UQ calculations can exploit the low-dimensional structure. When such structure is present, the resulting *dimension reduction* can enable otherwise infeasible UQ for expensive computer models with several inputs. For example, *sensitivity analysis* [5] may identify unimportant inputs (i.e., inputs whose variation does not affect the predictions), which may be fixed at nominal values for UQ. *Active subspaces* [6, 7] generalize this idea by identifying important linear combinations of variables; the active subspace's orthogonal complement contains directions in the input space along which a particular model output changes relatively little, on average. One can exploit these subspaces to reduce the dimension for UQ, and the reduced dimension is typically smaller than (and always bounded above by) the number of important inputs ranked by coordinate-based sensitivity analysis. However, estimating a model's active subspace requires access to the gradient (or an ap-

proximate gradient) of the input/output map. If the computer model does not produce gradients and its evaluations are too expensive for approximations (e.g., finite differences), then estimating active subspaces is challenging.

In statistical regression, related subspace-based dimension reduction exists under the moniker *sufficient dimension reduction* (SDR) [8]. Techniques include ordinary least squares (OLS) [9], sliced inverse regression (SIR) [10], sliced average variance estimation (SAVE) [11], and principal Hessian directions (pHd) [12]—among several others; see [8] for a comprehensive review. These techniques seek a low-dimensional subspace that is sufficient to characterize the relationship between predictors and response. The regression problem begins with predictor/response pairs—assumed to be independent draws from a joint predictor/response distribution—and constructs a *link function* that models the expected predictor/response relationship; for example, the response may be modeled as a multivariate polynomial of the predictors. SDR methods use the same predictor/response pairs to probe for a dimension reduction subspace without the gradient of the link function.

In this paper, we evaluate and analyze *inverse regression* methods, SIR and SAVE, as candidates for gradient-free subspace-based dimension reduction in deterministic computer models. To set the computer model appropriately, consider sampling the physical inputs according to a chosen design-of-experiments—e.g., a Latin hypercube design [13]. For each input sample, run the computer model and acquire an associated output quantity of interest. Then treat the set of input/output pairs as *data* for the inverse regression method. Such an idea is not without precedent. Cook used OLS, SIR, and pHd with appropriate variable transformations to identify important inputs and construct an efficient regression metamodel for a computer simulation of an environmental contaminant [8, Chapter 14]. Moreover, Cook's plots revealed simple relationships between the transformed inputs and the complex model's output [14, Section 4]. This is the first work we know of that applies SDR methods to computer models. More recently, Li, Lin, and Li [15] applied inverse regression tools for uncertainty quantification of computer models. However, neither paper carefully addresses the

mathematical subtleties involved in applying a technique developed for statistical regression to study a deterministic function. When applying SDR methods to deterministic computer models, the data set (i.e., a design-of-experiments in the physical inputs and associated computer model predictions) differs from the regression data set (independent draws of a joint predictor/response distribution). How does one interpret a dimension reduction subspace from SIR or SAVE when there is no random noise in the functional relationship between inputs and outputs? What does *sufficient dimension reduction* mean for deterministic functions?

To answer these questions, we interpret the inverse regression methods as methods for *ridge recovery* in deterministic functions. A function of several variables is called a *ridge function* when it can be written as another function of a few linear combinations of the variables [16]. Recent work by Fornasier, Schnass, and Vybiral [17] and related work by Tyagi and Cevher [18] attempt to recover the ridge directions (i.e., the weights of the linear combinations) from point evaluations of the function. Their approach exploits a connection between a linear measurement of a gradient vector and the function's directional derivative along the direction defined by the linear measurement. Constantine, Eftekhari, and Wakin [19] exploit a similar connection to estimate active subspaces without gradients. Ridge recovery is a deterministic analog of sufficient dimension reduction, so ridge recovery provides an appropriate mathematical setting in which to evaluate the inverse regression methods.

Our main results are the following. If we assume that the data supplied to SIR or SAVE come from point evaluations of a deterministic ridge function, then the SIR and SAVE subspaces are subspaces of the span of the ridge directions; this connects the SIR and SAVE algorithms to deterministic ridge recovery. We also interpret the elements of the matrices that define the SIR and SAVE subspaces as deterministic integrals on subsets of the input parameter space. This allows us to interpret SIR and SAVE as approximating integrals, which leads to new convergence analysis of the SIR and SAVE subspace estimates. We complement these results with numerical experiments on three deterministic ridge

functions: two quadratic functions and a simplified model of magnetohydrodynamics.

The remainder of the paper is structured as follows. Section 2 reviews (i) the essential theory of sufficient dimension reduction in statistical regression and (ii) the SIR and SAVE algorithms. Section 3 translates these results to deterministic ridge recovery and analyzes the convergence of the SIR and SAVE subspaces. Section 4 shows numerical examples that support the analyses using (i) two quadratic functions of 10 variables and (ii) a simplified model of magnetohydrodynamics with 5 input parameters. Section 5 contains concluding remarks. To improve readability, all proofs are in the Appendix.

2. Sufficient dimension reduction

We first review the essential theory of sufficient dimension reduction (SDR); our notation and development closely follow Cook's monograph *Regression Graphics: Ideas for Studying Regressions through Graphics* [8]. The theory of SDR provides a framework for subspace-based dimension reduction in statistical regression. A regression problem begins with predictor/response pairs $\{[\, y_i \,,\, \mathbf{x}_i^T \,]\}$, $i = 1, \ldots, N$, where $y_i \in \mathbb{R}$ and $\mathbf{x}_i \in \mathbb{R}^m$ denote the scalar-valued response and vector-valued predictors, respectively. These pairs are assumed to be independent realizations from the random vector $[\, y \,,\, \mathbf{x}^T \,]$ with (unknown) joint density $\pi(y, \mathbf{x})$. The object of interest in regression is the conditional random variable $y|\mathbf{x}$; a statistician uses the given predictor/response pairs to estimate statistics of $y|\mathbf{x}$—e.g., conditional moments or quantiles. SDR searches for a subspace of the predictors that is sufficient to describe $y|\mathbf{x}$ with statistics derived from the given predictor/response pairs.

2.1. Dimension reduction subspaces (DRS)

The basic tool of SDR is the *dimension reduction subspace* (DRS).

Definition 1 (Dimension reduction subspace [20])**.** *Consider a random vector* $[\, y\,, \mathbf{x}^T\,] \in \mathbb{R}^{m+1}$*, and let* $\boldsymbol{A} \in \mathbb{R}^{m\times n}$ *with* $n \leq m$ *be such that*

$$y \perp\!\!\!\perp \mathbf{x} | \boldsymbol{A}^T\mathbf{x}, \tag{1}$$

where $\perp\!\!\!\perp$ *denotes independence of* y *and* $\mathbf{x}$*. A* dimension reduction subspace $\mathcal{S}_{\text{DRS}}$ *for* $y|\mathbf{x}$ *is*

$$\mathcal{S}_{\text{DRS}} \;=\; \mathcal{S}_{\text{DRS}}(\boldsymbol{A}) \;=\; \text{colspan}(\boldsymbol{A}) \tag{2}$$

where $\text{colspan}(\boldsymbol{A})$ *denotes the column space of the* $\boldsymbol{A}$*.*

Equation (1) denotes the conditional independence of the response and predictors given $\boldsymbol{A}^T\mathbf{x}$. We can define a new random variable $y|\boldsymbol{A}^T\mathbf{x}$ that is the response conditioned on the n-dimensional vector $\boldsymbol{A}^T\mathbf{x}$. If $\mathcal{S}_{\text{DRS}}(\boldsymbol{A})$ is a DRS for $y|\mathbf{x}$, then the conditional CDF for $y|\mathbf{x}$ is the conditional CDF for $y|\boldsymbol{A}^T\mathbf{x}$,

$$F_{y|\mathbf{x}}(t) \;=\; F_{y|\boldsymbol{A}^T\mathbf{x}}(t), \qquad t \in \mathbb{R}. \tag{3}$$

See [8, Chapter 6]. If a regression admits a low-dimensional ($n < m$) DRS, then the predictor dimension can be reduced by considering $y|\boldsymbol{A}^T\mathbf{x}$.

The matrix $\boldsymbol{A}$ in (2) is not unique [21]. For any matrices $\boldsymbol{A}, \boldsymbol{B} \in \mathbb{R}^{m\times n}$ such that $\text{colspan}(\boldsymbol{A}) = \text{colspan}(\boldsymbol{B})$,

$$y \perp\!\!\!\perp \mathbf{x} | \boldsymbol{A}^T\mathbf{x} \quad \Leftrightarrow \quad y \perp\!\!\!\perp \mathbf{x} | \boldsymbol{B}^T\mathbf{x}. \tag{4}$$

Moreover, it is possible for two different subspaces to be DRSs for the same regression. Suppose a particular regression admits the DRS $\mathcal{S}_{\text{DRS}}(\boldsymbol{A})$ with $\boldsymbol{A} \in \mathbb{R}^{m\times n}$ and $n < m$, and let $\mathbf{v} \in \mathbb{R}^m$ such that $\mathbf{v} \notin \text{colspan}(\boldsymbol{A})$. Consider the $m \times (n+1)$ matrix $\boldsymbol{B} = \begin{bmatrix}\boldsymbol{A} & \mathbf{v}\end{bmatrix}$. The subspace $\text{colspan}(\boldsymbol{B})$ is also a DRS for the same regression. However, $\mathcal{S}_{\text{DRS}}(\boldsymbol{A})$ is the preferred subspace since $\dim\left(\mathcal{S}_{\text{DRS}}(\boldsymbol{A})\right) \leq \dim(\mathcal{S}_{\text{DRS}}(\boldsymbol{B}))$ where $\dim(\cdot)$ denotes the subspace dimension. It is not necessary that a DRS reduce the predictor dimension of the regression. In fact, all regressions admit the trivial DRS, $\mathbb{R}^m$ [8, Chapter 6]. Since the goal of SDR is to reduce the predictor dimension, we are interested in finding the one which has minimum dimension. The *minimum DRS* maximizes the dimension reduction for a particular regression.

Definition 2 (Minimum DRS [20]). *A subspace $\mathcal{S}$ is called a* minimum DRS *for $y|\mathbf{x}$ if it is a DRS and $\dim(\mathcal{S}) \leq \dim(\mathcal{S}_{\text{DRS}})$, where $\mathcal{S}_{\text{DRS}}$ is any DRS for $y|\mathbf{x}$.*

Since the predictor dimension is finite, a minimum DRS always exists—though the minimum DRS may be $\mathbb{R}^m$. Existence of the minimum DRS indicates that all regression problems exhibit a minimum dimension for sufficient dimension reduction. This minimum dimension, referred to as the *structural dimension* of $y|\mathbf{x}$, represents the largest possible reduction of the predictor space while maintaining the structure of $y|\mathbf{x}$ [20]. Despite uniqueness of the structural dimension, a unique minimum DRS is not guaranteed—as the following example from Cook [21] illustrates.

Example 1. *Assume that $\mathbf{x} \in \mathbb{R}^2$ has a uniform marginal density on the unit circle. Suppose that the random variable $y|\mathbf{x} = x_1^2 + \epsilon$, where ϵ is random noise that is independent of $\mathbf{x}$. This implies that $\mathcal{S}_{\text{DRS}}([\,1\,,\,0\,]^T)$ is a DRS for this problem. However, the marginal density of $\mathbf{x}$ is such that $1 = x_1^2 + x_2^2$. This means $y|\mathbf{x} = 1 - x_2^2 + \epsilon$. From this expression, $\mathcal{S}_{\text{DRS}}([\,0\,,\,1\,]^T)$ is also a DRS. Both of these DRSs have dimension 1 and there does not exist a DRS of dimension 0 (i.e., $\{\mathbf{0}\}$ is not a DRS). Therefore, there does not exist a unique minimum DRS for this regression.*

We wish to find a uniquely defined minimum DRS to guarantee the well-posedness of the SDR. To this end, we consider the *central DRS*.

Definition 3 (Central DRS [20]). *A subspace $\mathcal{S}_{y|\mathbf{x}}$ is a* central DRS *for $y|\mathbf{x}$ if it is a DRS and $\mathcal{S}_{y|\mathbf{x}} \subseteq \mathcal{S}_{\text{DRS}}$, where $\mathcal{S}_{\text{DRS}}$ is any DRS for $y|\mathbf{x}$.*

When the central DRS exists it will be the intersection of all other DRSs,

$$\mathcal{S}_{y|\mathbf{x}} \;=\; \bigcap \mathcal{S}_{\text{DRS}} \tag{5}$$

for all $\mathcal{S}_{\text{DRS}}$ for $y|\mathbf{x}$. The intersection in (5) always defines a subspace, but this subspace need not satisfy the conditional independence from Definition 1. Therefore, a regression need not admit a central DRS—see, e.g., example 1.

However, when a central DRS exists, it is the unique minimum DRS for $y|\mathbf{x}$ [8, Chapter 6]. For this reason, finding the central DRS is the goal of SDR—provided that the central DRS exists.

Existence of the central DRS can be proved under particular conditions; one case is via *location regressions*. A location regression is a regression for which $y \perp\!\!\!\perp \mathbf{x} | \mathbb{E}[y|\mathbf{x}]$ [8, Chapter 6].

Theorem 1 ([8]). *Suppose that $\mathcal{S}_1$ and $\mathcal{S}_2$ are two DRSs for a location regression $y|\mathbf{x}$ such that $y \perp\!\!\!\perp \mathbf{x} | \mathbb{E}[y|\mathbf{x}]$. Additionally, assume that $\mathbb{E}[y|\mathbf{x}]$ may be expressed as a convergent power series for all $\mathbf{x}$. Then, $\mathcal{S}_1 \cap \mathcal{S}_2$ is also a DRS.*

Recall that the central DRS is the intersection of all DRSs provided that this intersection is a DRS as well. Theorem 1 provides conditions under which the intersection of any two DRSs is a DRS which implies that a central DRS exists. A second case uses *M-sets* to guarantee existence of the central DRS.

Definition 4 (M-set [22]). *A set $\mathcal{A} \in \mathbb{R}^p \times \mathbb{R}^q$ is called an* M-set *if, for any two elements $(\mathbf{u}_1, \mathbf{v}_1), (\mathbf{u}_2, \mathbf{v}_2) \in \mathcal{A}$, there is a collection of pairs $\{(\mathbf{u}^{(i)}, \mathbf{v}^{(i)})\}_{i=1}^{M}$ in $\mathcal{A}$ such that $(\mathbf{u}^{(1)}, \mathbf{v}^{(1)}) = (\mathbf{u}_1, \mathbf{v}_1)$, $((\mathbf{u}^{(M)}, \mathbf{v}^{(M)}) = (\mathbf{u}_2, \mathbf{v}_2)$, and either $\mathbf{u}^{(i)} = \mathbf{u}^{(i+1)}$ or $\mathbf{v}^{(i)} = \mathbf{v}^{(i+1)}$ for $i = 1, \ldots, M-1$.*

Using this definition, the following theorem provides conditions for the existence of a central DRS.

Theorem 2 ([22]). *Suppose that $\mathcal{S}_1(\boldsymbol{A}_1)$ and $\mathcal{S}_2(\boldsymbol{A}_2)$ are two DRSs for $y|\mathbf{x}$ where $\boldsymbol{A}_1 \in \mathbb{R}^{m \times n_1}$ and $\boldsymbol{A}_2 \in \mathbb{R}^{m \times n_2}$. Let $\mathcal{S}_3(\boldsymbol{A}_3) = \mathcal{S}_1(\boldsymbol{A}_1) \cap \mathcal{S}_2(\boldsymbol{A}_2)$. If, for each $\mathbf{u} \in \operatorname{colspan}(\boldsymbol{A}_3)$, the set*

$$\left\{ (\boldsymbol{A}_1^T \mathbf{x}, \boldsymbol{A}_2^T \mathbf{x}) \in \mathbb{R}^{n_1} \times \mathbb{R}^{n_2} \,\middle|\, \mathbf{u} = \boldsymbol{A}_3^T \mathbf{x} \right\} \tag{6}$$

is an M-set in $\mathbb{R}^{n_1} \times \mathbb{R}^{n_2}$, then $\mathcal{S}_1 \cap \mathcal{S}_2$ is also a DRS.

Any connected set is an M-set. However, a set need not be connected to be an M-set. This provides the set of existence conditions that are the most applicable to the next section. These conditions place restrictions on the marginal density of the predictors $\pi_{\mathbf{x}}(\mathbf{x})$.

Theorem 3 ([8]). *Suppose that $\mathcal{S}_1$ and $\mathcal{S}_2$ are two DRSs for $y|\mathbf{x}$ where the marginal density of the predictors $\pi_{\mathbf{x}}(\mathbf{x})$ has support over a convex set. Then, $\mathcal{S}_1 \cap \mathcal{S}_2$ is also a DRS.*

This existence condition is relatively straightforward to interpret, but such knowledge of the predictor space is not generally assumed in regression. However, this condition is useful for translating SDR theory to deterministic functions in section 3. For the remainder of this discussion, we assume that the regression is such that a central DRS exists.

The central DRS has two properties that are important in the development of the inverse regression methods in the next sections. The first is the transformation of the central DRS under affine transformations of the predictors.

Property 1 ([20]). *Let $\mathbf{z} = \boldsymbol{B}\mathbf{x}+\mathbf{b}$ be an affine transformation of the predictors $\mathbf{x}$, where $\boldsymbol{B} \in \mathbb{R}^{m\times m}$ is full rank and $\mathbf{b} \in \mathbb{R}^m$. Additionally, let $\mathcal{S}_{y|\mathbf{x}}$ be the central DRS for $y|\mathbf{x}$ with basis matrix $\boldsymbol{A}$. Then,* $\operatorname{colspan}((\boldsymbol{B}^{-1})^T\boldsymbol{A})$ *is the central DRS for $y|\mathbf{z}$.*

Property 1 allows for standardization of the predictors as follows. Let $\boldsymbol{\mu} = \mathbb{E}[\mathbf{x}]$ and $\boldsymbol{\Sigma}_{\mathbf{x}} = \operatorname{Cov}[\mathbf{x}]$, and suppose $\boldsymbol{\Sigma}_{\mathbf{x}}$ is full rank with Cholesky factorization $\boldsymbol{\Sigma}_{\mathbf{x}} = \boldsymbol{L}\boldsymbol{L}^T$. Define the standardized predictor $\mathbf{z} = \boldsymbol{L}^{-1}(\mathbf{x}-\boldsymbol{\mu})$ so that $\mathbb{E}[\mathbf{z}] = \mathbf{0}$ and $\operatorname{Cov}[\mathbf{z}] = \boldsymbol{I}$. The central DRS for $\mathbf{z}$ relates to the central DRS for $\mathbf{x}$ by (with an abuse of notation) $\mathcal{S}_{y|\mathbf{x}} = \boldsymbol{L}^{-T}\mathcal{S}_{y|\mathbf{z}}$. This leads to the following assumption.

Assumption 1 (Standardized predictors). *Assume that the predictors $\mathbf{x}$ are standardized so that $\mathbb{E}[\mathbf{x}] = \mathbf{0}$ and $\operatorname{Cov}[\mathbf{x}] = \boldsymbol{I}$.*

For the remainder of this discussion we assume that Assumption 1 holds unless otherwise stated.

The second property of the central DRS is the containment of the subspace under transformation of the response.

Property 2 ([23]). *Let $\mathcal{S}_{y|\mathbf{x}}$ be the central DRS for $y|\mathbf{x}$ and $h(y)$ be any function of the response. Then,*

$$\mathcal{S}_{h(y)|\mathbf{x}} \subseteq \mathcal{S}_{y|\mathbf{x}} \tag{7}$$

and $\mathcal{S}_{h(y)|\mathbf{x}} = \mathcal{S}_{y|\mathbf{x}}$ when $h(y)$ is strictly monotonic.

Property 2 states the central DRS for $h(y)|\mathbf{x}$ is contained within the central DRS for $y|\mathbf{x}$. The following example from [8, Chapter 6] shows that these subspaces need not always be equal.

Example 2. *Let $y|\mathbf{x}$ be normally distributed with $\mathbb{E}[y|\mathbf{x}] = 0$ and $\mathrm{Var}[y|\mathbf{x}] = \mathrm{Var}\left[y|\boldsymbol{A}^T\mathbf{x}\right]$ where $\boldsymbol{A} \in \mathbb{R}^{m\times n}$ contains the basis for the central DRS for $y|\mathbf{x}$. Assume that $h(y)$ maps negative values of y to $h_1 \in \mathbb{R}$ and non-negative values of y to $h_2 \in \mathbb{R}$. Then, $\mathbb{P}[h(y) = h_1|\mathbf{x}] = \mathbb{P}[h(y) = h_1] = 1/2$ for all $\mathbf{x}$, which implies that $\mathcal{S}_{h(y)|\mathbf{x}} = \{\mathbf{0}\}$. However, $\mathcal{S}_{y|\mathbf{x}} \neq \{\mathbf{0}\}$ in general. Thus, $\mathcal{S}_{h(y)|\mathbf{x}}$ is a strict subset of $\mathcal{S}_{y|\mathbf{x}}$.*

The goal of SDR is to estimate the central DRS for the regression from the given response/predictor pairs. We formalize the *SDR problem* below.

Problem 1 (SDR problem). *Given response/predictor pairs $\{[\,y_i\,,\,\mathbf{x}_i^T\,]\}$, with $i = 1,\ldots,N$, assumed to be independent draws from a random vector $[\,y\,,\,\mathbf{x}^T\,]$ with joint density $\pi(y,\mathbf{x})$, compute a basis $\boldsymbol{A} \in \mathbb{R}^{m\times n}$ for the central DRS $\mathcal{S}_{y|\mathbf{x}}$ of the random variable $y|\mathbf{x}$.*

The ideas discussed in this section provide the theoretical framework for the algorithmic approaches for Problem 1. Next, we introduce two algorithms for SDR—sliced inverse regression (SIR) and sliced average variance estimation (SAVE)—which belong to the class of *inverse regression* methods. These methods use statistical characteristics of the inverse regression $\mathbf{x}|y$ to approximate the central DRS of $y|\mathbf{x}$. The inverse regression is treated as m univariate random variables conditioned on the scalar-valued response which avoids dealing with the high-dimensional forward regression $y|\mathbf{x}$. SIR and SAVE construct subspaces from $\mathbf{x}|y$ which provide insight into the central DRS $\mathcal{S}_{y|\mathbf{x}}$.

2.2. Sliced inverse regression

Sliced inverse regression approximates the central DRS using the expectation of the inverse regression $\mathbb{E}[\mathbf{x}|y]$, referred to as the *inverse regression function*. This function draws a curve through the m-dimensional predictor space parameterized by the scalar-valued response. Let $\mathcal{S}_{\mathbb{E}[\mathbf{x}|y]}$ denote the subspace containing the inverse regression function as it varies over all values of y in the support of $\pi(y, \mathbf{x})$. The goal of SIR is to estimate this subspace using the response/predictor pairs.

Before providing the algorithm for SIR, we discuss the following condition.

Condition 1 (Linearity condition [10])**.** *Let $\mathcal{S}_{y|\mathbf{x}}(\boldsymbol{A})$ be the central DRS for $y|\mathbf{x}$. The linearity condition is satisfied if for any $\mathbf{b} \in \mathbb{R}^m$ the conditional expectation $\mathbb{E}\left[\mathbf{b}^T\mathbf{x}|\boldsymbol{A}^T\mathbf{x}\right]$ is linear in $\boldsymbol{A}^T\mathbf{x}$. That is,*

$$\mathbb{E}\left[\mathbf{b}^T\mathbf{x}|\boldsymbol{A}^T\mathbf{x}\right] = c_0 + c_1\mathbf{a}_1^T\mathbf{x} + \cdots + c_n\mathbf{a}_n^T\mathbf{x} \tag{8}$$

for some constants $c_0, c_1, \ldots, c_n$ where $\mathbf{a}_i$ denotes the i^{th} column of $\boldsymbol{A}$.

The linearity condition places a restriction on the expectation of the predictors $\mathbf{x}$ conditioned on $\boldsymbol{A}^T\mathbf{x}$, where $\boldsymbol{A}$ contains a basis for the central DRS. However, the goal of Problem 1 is to compute a basis for the central DRS. This implies that we cannot check Condition 1 prior to solving the SDR problem since such an $\boldsymbol{A}$ would be unknown. Therefore, Condition 1 is generally extended to all matrices that generate a DRS for the regression. This extended condition (and Condition 1) is satisfied when the marginal density $\pi_{\mathbf{x}}(\mathbf{x})$ is elliptically symmetric [24]. One well-known example of an elliptically symmetric distribution is the multivariate standard Gaussian. In cases where elliptic symmetry of the predictors cannot be verfied, we may still be able to assume Condition 1 to be nearly satisfied. Hall and Li [25] have shown that for increasing dimensions of $\mathbf{x}$, $\mathbb{E}\left[\mathbf{b}^T\mathbf{x}|\boldsymbol{A}^T\mathbf{x}\right]$ approaches linearity under certain conditions. For the remainder, we assume that the linearity condition holds.

Algorithm 1 estimates the subspace containing the inverse regression function $\mathcal{S}_{\mathbb{E}[\mathbf{x}|y]}$. The following theorem provides the relationship between $\mathcal{S}_{\mathbb{E}[\mathbf{x}|y]}$ and the central DRS for $y|\mathbf{x}$.

Algorithm 1 Sliced inverse regression [10]

Given: N samples $\{[\, y_i \,,\, \mathbf{x}_i^T]\}$, $i = 1, \ldots, N$, drawn independently according to $\pi(y, \mathbf{x})$.

Assumptions: Assumption 1 holds and Condition 1 is satisfied.

1. Define a partition of the response space,

$$y_{\min} = \tilde{y}_0 < \tilde{y}_1 < \cdots < \tilde{y}_{R-1} < \tilde{y}_R = y_{\max}, \tag{9}$$

   where $y_{\min} = \min_i y_i$ and $y_{\max} = \max_i y_i$. For $r = 1, \ldots, R$, let $J_r = [\tilde{y}_{r-1}, \tilde{y}_r]$ denote the rth slice of the partition, and let $\mathcal{I}_r \subset \{1, \ldots, N\}$ be the set of indices i for which $y_i \in J_r$.

2. For $r = 1, \ldots, R$, compute the sample mean $\hat{\boldsymbol{\mu}}(r)$ of the predictors whose associated responses are in J_r,

$$\hat{\boldsymbol{\mu}}(r) = \frac{1}{N_r} \sum_{i \in \mathcal{I}_r} \mathbf{x}_i \tag{10}$$

   where N_r is the cardinality of $\mathcal{I}_r$.

3. Compute the weighted sample covariance matrix

$$\hat{\boldsymbol{D}}_{\text{SIR}} \;=\; \frac{1}{N} \sum_{r=1}^{R} N_r \, \hat{\boldsymbol{\mu}}(r) \hat{\boldsymbol{\mu}}(r)^T. \tag{11}$$

4. Compute the eigendecomposition,

$$\hat{\boldsymbol{D}}_{\text{SIR}} \;=\; \hat{\boldsymbol{W}} \hat{\boldsymbol{\Lambda}} \hat{\boldsymbol{W}}^T, \tag{12}$$

   where the eigenvalues are in descending order $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \cdots \geq \hat{\lambda}_m \geq 0$ and the eigenvectors are orthonormal.

5. Let $\hat{\boldsymbol{A}} \in \mathbb{R}^{m \times n}$ be the first n eigenvectors of $\hat{\boldsymbol{D}}_{\text{SIR}}$.

Theorem 4 ([8])**.** *Assume that Condition 1 is satisfied. Then,*

$$\mathcal{S}_{\mathbb{E}[\mathbf{x}|y]} \subseteq \mathcal{S}_{y|\mathbf{x}}. \tag{13}$$

Theorem 4 states that, under the linearity condition, the curve drawn out by the inverse regression function is contained within the central DRS. This implies that the curve does not vary in directions orthogonal to the central DRS so that the covariance matrix of the inverse regression function is degenerate along these directions. Cook formalizes this idea by showing that $\mathcal{S}_{\mathbb{E}[\mathbf{x}|y]} = \mathcal{S}_{\mathrm{Cov}[\mathbb{E}[\mathbf{x}|y]]}$ where $\mathcal{S}_{\mathrm{Cov}[\mathbb{E}[\mathbf{x}|y]]} = \mathrm{colspan}(\mathrm{Cov}\left[\mathbb{E}\left[\mathbf{x}|y\right]\right])$ [8, Chapter 11]. Therefore, knowledge of the subspace $\mathcal{S}_{\mathrm{Cov}[\mathbb{E}[\mathbf{x}|y]]}$ provides at least part of $\mathcal{S}_{y|\mathbf{x}}$.

Algorithm 1 constructs estimates of the covariance matrix $\mathrm{Cov}\left[\mathbb{E}\left[\mathbf{x}|y\right]\right]$ using a partitioning of the observed response range. This *slicing* approach is motivated by Property 2 of the central DRS. Let $h : \mathbb{R} \to \{1, \ldots, N\}$ map the response to a slice of the response space,

$$y \in J_r \Longrightarrow h(y) = r, \tag{14}$$

where J_r denotes the rth slice of the observed range from (9) in Algorithm 1. Without loss of generality, we assume the slices are constructed such that $\mathbb{P}[y \in J_r] > 0$ for all J_r. If this is not true for some slice, then that slice may be combined with either of the neighboring slices without altering the behavior of the function h.

Combining the slice map h with Theorem 4 yields

$$\mathcal{S}_{\mathrm{Cov}[\mathbb{E}[\mathbf{x}|h(y)]]} = \mathcal{S}_{\mathbb{E}[\mathbf{x}|h(y)]} \subseteq \mathcal{S}_{h(y)|\mathbf{x}} \subseteq \mathcal{S}_{y|\mathbf{x}}. \tag{15}$$

The matrix $\mathrm{Cov}\left[\mathbb{E}\left[\mathbf{x}|h(y)\right]\right]$ is known as the *SIR matrix*, and its dominant n-dimensional eigenspace is the n-dimensional *SIR subspace*. Equation (15) motivates the SIR algorithm, which constructs the approximation $\hat{\boldsymbol{D}}_{\mathrm{SIR}} \approx \mathrm{Cov}\left[\mathbb{E}\left[\mathbf{x}|h(y)\right]\right]$. The n-dimensional eigenspace of $\hat{\boldsymbol{D}}_{\mathrm{SIR}}$ estimates the dominant n-dimensional SIR subspace.

Eigenvectors of $\mathrm{Cov}\left[\mathbb{E}\left[\mathbf{x}|h(y)\right]\right]$ associated with nonzero eigenvalues provide a basis for the SIR subspace. However, eigenvalues of the approximated SIR

matrix $\hat{\boldsymbol{D}}_{\text{SIR}}$ are unlikely to be exactly zero, which raises the issue of determining the dimensionality of the SIR subspace from the eigenvalue estimates. Li [10] and Cook [11] propose significance tests based on the distribution of the average of the $m - n$ trailing estimated eigenvalues. These testing methods extend to the SAVE algorithm in section 2.3.

For a fixed number of slices, the SIR algorithm has been shown to be $N^{-1/2}$ consistent in approximating $\mathcal{S}_{\text{Cov}[\mathbb{E}[\mathbf{x}|h(y)]]}$ [10]. In principle, increasing the number of slices may provide improved estimation of the central DRS as $h(y)$ approaches strict monotonicity. However, in practice, the success of Algorithm 1 has been found to be relatively insensitive to the number of slices. The number of slices should be chosen such that sufficient samples are present in each slice to estimate the conditional expectation accurately. For this reason, Li [10] suggests constructing slices such that the response samples are distributed nearly equally.

The popularity of SIR is due in part to the straightforward-to-implement Algorithm 1. However, SIR fails when the inverse regression is symmetric about $\mathbf{x} = \mathbf{0}$ for all values of the response [11]. This failure occurs because $\mathbb{E}[\mathbf{x}|y] = \mathbf{0}$ for all y. The subspace is then $\mathcal{S}_{\mathbb{E}[\mathbf{x}|y]} = \{\mathbf{0}\}$, which still satisfies the containment in Theorem 4. However, recovering the origin does not help discover the central DRS. This failure of SIR led to alternate methods, such as SAVE, which use higher moments of the inverse regression.

2.3. Sliced average variance estimation

Where SIR uses the expectation of the inverse regression, sliced average variance estimation (SAVE) uses the variance. Li [10] recognized the potential for $\text{Cov}[\mathbf{x}|y]$ to provide insights into the central DRS by noting that, under Assumption 1,

$$\mathbb{E}[\text{Cov}[\mathbf{x}|y]] = \text{Cov}[\mathbf{x}] - \text{Cov}[\mathbb{E}[\mathbf{x}|y]] = \boldsymbol{I} - \text{Cov}[\mathbb{E}[\mathbf{x}|y]]. \tag{16}$$

Equation (16) implies that $\mathbb{E}[\boldsymbol{I} - \text{Cov}[\mathbf{x}|y]]$ may be useful in estimating $\mathcal{S}_{y|\mathbf{x}}$. Cook suggests using $\mathbb{E}\left[(\boldsymbol{I} - \text{Cov}[\mathbf{x}|y])^2\right]$, which has positive eigenvalues [11].

The SAVE algorithm attempts to estimate the column space of this matrix, which we denote by $\mathcal{S}_{\mathbb{E}[(\boldsymbol{I}-\mathrm{Cov}[\mathbf{x}|y])^2]}$. Before detailing the SAVE algorithm, we introduce a condition that allows one to relate the subspace SAVE produces to the central DRS.

Condition 2 (Constant covariance condition [23])**.** *Let* $\mathcal{S}_{y|\mathbf{x}}(\boldsymbol{A})$ *be the central DRS for* $y|\mathbf{x}$*. The* constant covariance condition *is satisfied if the conditional covariance matrix* $\mathrm{Cov}\left[\mathbf{x}|\boldsymbol{A}^T\mathbf{x}\right]$ *is a constant matrix.*

Similar to Condition 1, confirming the constant covariance condition requires knowledge of the central DRS. Therefore, we consider an extension that requires $\mathrm{Cov}\left[\mathbf{x}|\boldsymbol{A}^T\mathbf{x}\right]$ to be constant for all $\boldsymbol{A}$ that generate a DRS for $y|\mathbf{x}$. Elliptic symmetry of the predictors implies this extension [24].

Algorithm 2 computes an estimate for $\mathcal{S}_{\mathbb{E}[(\boldsymbol{I}-\mathrm{Cov}[\mathbf{x}|y])^2]}$. The following theorem shows why this subspace is useful in discovering the central DRS.

Theorem 5 ([23])**.** *Assume that Conditions 1 and 2 are satisfied. Then,*

$$\mathcal{S}_{\mathbb{E}[(\boldsymbol{I}-\mathrm{Cov}[\mathbf{x}|y])^2]} \subseteq \mathcal{S}_{y|\mathbf{x}}. \tag{22}$$

Theorem 5 implies that the subspace $\mathcal{S}_{\mathbb{E}[(\boldsymbol{I}-\mathrm{Cov}[\mathbf{x}|y])^2]}$ reveals at least part of the central DRS. Applying the slicing map h from (14) to Theorem 5,

$$\mathcal{S}_{\mathbb{E}[(\boldsymbol{I}-\mathrm{Cov}[\mathbf{x}|h(y)])^2]} \subseteq \mathcal{S}_{h(y)|\mathbf{x}} \subseteq \mathcal{S}_{y|\mathbf{x}}. \tag{23}$$

The matrix $\mathbb{E}\left[(\boldsymbol{I}-\mathrm{Cov}\left[\mathbf{x}|h(y)\right])^2\right]$ is the *SAVE matrix*, and its n-dimensional eigenspace is the n-dimensional SAVE subspace. Equation (23) motivates the SAVE algorithm for estimating the central DRS first developed by Cook and Weisberg [11]. The SAVE algorithm computes the numerical approximation $\hat{\boldsymbol{D}}_{\mathrm{SAVE}} \approx \mathbb{E}\left[(\boldsymbol{I}-\mathrm{Cov}\left[\mathbf{x}|h(y)\right])^2\right]$. The dominant n-dimensional eigenspace of $\hat{\boldsymbol{D}}_{\mathrm{SAVE}}$ estimates the n-dimensional SAVE subspace.

The SAVE algorithm approximates the subspace $\mathcal{S}_{\mathbb{E}[(\boldsymbol{I}-\mathrm{Cov}[\mathbf{x}|y])^2]}$ with $N^{-1/2}$ consistency [23]. Increasing the number of slices improves the estimate but suffers the same drawbacks as in SIR. In practice, SAVE has been shown to perform poorly compared to SIR when fewer samples are available. This is due

Algorithm 2 Sliced average variance estimation [23]

Given: N samples $\{[\, y_i \,,\, \mathbf{x}_i^T]\}$, $i = 1, \ldots, N$, drawn independently according to $\pi(y, \mathbf{x})$.

Assumptions: Assumption 1 holds and Conditions 1 and 2 are satisfied.

1. Define a partition of the response space,

$$y_{\min} = \tilde{y}_0 < \tilde{y}_1 < \cdots < \tilde{y}_{R-1} < \tilde{y}_R = y_{\max}, \tag{17}$$

   where $y_{\min} = \min_i y_i$ and $y_{\max} = \max_i y_i$. For $r = 1, \ldots, R$, let $J_r = [\tilde{y}_{r-1}, \tilde{y}_r]$ denote the rth slice of the partition, and let $\mathcal{I}_r \subset \{1, \ldots, N\}$ be the set of indices i for which $y_i \in J_r$.

2. For $r = 1, \ldots, R$,

   (a) Compute the sample mean $\hat{\boldsymbol{\mu}}(r)$ of the predictors whose associated responses are in the J_r,

$$\hat{\boldsymbol{\mu}}(r) = \frac{1}{N_r} \sum_{i \in \mathcal{I}_r} \mathbf{x}_i \tag{18}$$

   where N_r is the cardinality of $\mathcal{I}_r$.

   (b) Compute the sample covariance $\hat{\boldsymbol{\Sigma}}(r)$ of the predictors whose associated responses are in J_r,

$$\hat{\boldsymbol{\Sigma}}(r) = \frac{1}{N_r - 1} \sum_{i \in \mathcal{I}_r} \left(\mathbf{x}_i - \hat{\boldsymbol{\mu}}(r)\right) \left(\mathbf{x}_i - \hat{\boldsymbol{\mu}}(r)\right)^T \tag{19}$$

3. Compute the matrix,

$$\hat{\boldsymbol{D}}_{\text{SAVE}} \;=\; \frac{1}{N} \sum_{r=1}^{R} N_r \left(\boldsymbol{I} - \hat{\boldsymbol{\Sigma}}(r)\right)^2. \tag{20}$$

4. Compute the eigendecomposition,

$$\hat{\boldsymbol{D}}_{\text{SAVE}} \;=\; \hat{\boldsymbol{W}} \hat{\boldsymbol{\Lambda}} \hat{\boldsymbol{W}}^T, \tag{21}$$

   where the eigenvalues are in descending order $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \cdots \geq \hat{\lambda}_m \geq 0$ and the eigenvectors are orthonormal.

5. Let $\hat{\boldsymbol{A}} \in \mathbb{R}^{m \times n}$ be the first n eigenvectors of $\hat{\boldsymbol{D}}_{\text{SAVE}}$.

to difficulties approximating the covariance within the slices with few samples. For this reason, Cook [26] suggests that both methods be used when attempting to approximate the central DRS.

The SDR theory and the SIR and SAVE algorithms apply to regressions, where the predictor/response pairs are assumed to be independent draws from a joint distribution. Example applications include studies on mussel sizes and public school test scores [8][Chapter 11]. In these examples, randomness in responses model inaccurate measurements of the mussels or unknown causes of test score variation. However, not all problems exhibit this type of randomness. In particular, computer simulations are better modeled by a deterministic relationship between y and $\mathbf{x}$. In what follows, we interpret SDR—specifically, the approximation techniques SIR and SAVE—in the context of deterministic functions.

3. Ridge recovery

A physical model is a map from the space of physical inputs to the space of predictions. We can represent a physical model by a deterministic function, $y = f(\mathbf{x})$, where $\mathbf{x}$ is an m-vector of input variables and y is a scalar-valued quantity of interest. We assume the physical model includes a characterization of the reasonable or feasible set of model inputs; this characterization is quantified by a probability density function $\rho(\mathbf{x})$. This differs from regression, where the relationship between y and $\mathbf{x}$ contains randomness; i.e., many y's are associated with the same $\mathbf{x}$. Furthermore, the input density is not known a priori in regression. The marginal density $\pi_{\mathbf{x}}(\mathbf{x})$ in regression is comparable to $\rho(\mathbf{x})$ in the sense that both are probability density functions on the space of $\mathbf{x}$. But the roles and accessibility of ρ and $\pi_{\mathbf{x}}$ are key differences between regression and physical modeling.

The notion of dimension reduction in deterministic functions that most closely resembles SDR in regression is found in *ridge functions*; we follow Pinkus' development of the basics of ridge functions [16].

Definition 5 (Ridge function [16]). *A function $f : \mathbb{R}^m \to \mathbb{R}$ is a* ridge function *if there exists matrix $\boldsymbol{A} \in \mathbb{R}^{m\times n}$, with $n \leq m$, and a function $g : \mathbb{R}^n \to \mathbb{R}$ such that*

$$y \;=\; f(\mathbf{x}) \;=\; g(\boldsymbol{A}^T\mathbf{x}). \tag{24}$$

The columns of $\boldsymbol{A}$ are the directions *of the ridge function.*

It is convenient to introduce additional nomenclature for the ridge function components to enable careful comparison with SDR in regression. We call $\boldsymbol{A}$ the *direction matrix*. The column space of the direction matrix is the *ridge subspace*, which we denote by $\mathcal{S}_{\mathrm{R}}(\boldsymbol{A})$. We call g the *ridge profile* of f. A ridge function is low-dimensional in the sense that f is constant along directions in the space of $\mathbf{x}$ that are orthogonal to the ridge subspace. That is, if $\boldsymbol{A}^T\mathbf{b} = \mathbf{0}$ for some $\mathbf{b} \in \mathbb{R}^m$, then

$$f(\mathbf{x}+\mathbf{b}) = g(\boldsymbol{A}^T(\mathbf{x}+\mathbf{b})) = g(\boldsymbol{A}^T\mathbf{x}) = f(\mathbf{x}). \tag{25}$$

For thorough comparison with SDR, we append the following two extreme cases to the definition of a ridge function: i) $\boldsymbol{A} = \mathbf{0}$ and ii) $n = m$. In SDR (see section 2), $\boldsymbol{A} = \mathbf{0}$ implies $y \perp\!\!\!\perp \mathbf{x}$, which means that the predictors have no influence on the random response. In the deterministic function case, $y = f(\mathbf{x}) = g(\boldsymbol{A}^T\mathbf{x}) = g(\mathbf{0})$ for any $\mathbf{x} \in \mathbb{R}^m$. This implies that y does not change with changing $\mathbf{x}$, so y is a constant. At the other extreme, $n = m$ does not result in dimension reduction. This case can be interpreted as the deterministic analog of the trivial DRS.

Recall from Problem 1 that the SDR problem searches for dimension reduction in regression problems from a given data set of predictor/response pairs. We seek to understand how the theory of SDR should be interpreted when applied to a similar data set constructed from a deterministic function. To that end, we seek to connect the DRS to the ridge subspace. The next theorem provides insight into this connection.

Theorem 6. *Let (Ω, Σ, P) be a probability triple with sample space Ω, σ-algebra Σ, and probability measure P. Suppose that $\mathbf{x} : \Omega \to \mathbb{R}^m$ and $y : \Omega \to \mathbb{R}$ are*

random variables related by a measurable function $f : \mathbb{R}^m \to \mathbb{R}$ *so that* $y = f(\mathbf{x})$. *Let* $\boldsymbol{A} \in \mathbb{R}^{m\times n}$ *be a constant matrix. Then* $y \perp\!\!\!\perp \mathbf{x}|\boldsymbol{A}^T\mathbf{x}$ *if and only if* $y = g(\boldsymbol{A}^T\mathbf{x})$ *where* $g : \mathbb{R}^n \to \mathbb{R}$ *is a measurable function.*

The proof of Theorem 6 is in section Appendix A.2. This theorem states that conditional independence of the inputs and output for a deterministic function is equivalent to the function being a ridge function. Hence, the ridge subspace may be interpreted as a DRS for deterministic functions.

For the remainder of this section, we translate the properties of the DRS in section 2.1 to the deterministic case; we begin with the direction matrix $\boldsymbol{A}$ in (30). Similar to the DRS from Definition 1, the matrix defining the ridge subspace is not unique; what matters is its column space. The following lemma formalizes this; a proof is in section Appendix A.3. Let $y = f(\mathbf{x})$ with $\mathbf{x} \in \mathbb{R}^m$ and $y \in \mathbb{R}$. Suppose that $f(\mathbf{x})$ may be written $y = f(\mathbf{x}) = g_1(\boldsymbol{A}_1^T\mathbf{x})$ for some $\boldsymbol{A}_1 \in \mathbb{R}^{m\times n_1}$ and $g_1 : \mathbb{R}^{n_1} \to \mathbb{R}$. If $\boldsymbol{A}_2 \in \mathbb{R}^{m\times n_2}$ is such that $\text{colspan}(\boldsymbol{A}_1) = \text{colspan}(\boldsymbol{A}_2)$, then there exists $g_2 : \mathbb{R}^{n_2} \to \mathbb{R}$ such that

$$y \;=\; f(\mathbf{x}) \;=\; g_2(\boldsymbol{A}_2^T\mathbf{x}). \tag{26}$$

Lemma 3 makes two important points. First, the directions of the ridge function are not unique in defining the ridge subspace, so we focus the analysis on the subspace; we denote a ridge subspace by $\mathcal{S}_{\text{R}}$ when the direction matrix is not relevant. Second, the function that links the transformed inputs to the quantity of interest depends on the direction matrix. This highlights the fact that the values $\boldsymbol{A}_1^T\mathbf{x}$ or $\boldsymbol{A}_2^T\mathbf{x}$ are inputs to a specific deterministic function. This differs from regression, where the transformed predictors are not tied to a function but are instead associated with a subspace that maintains the structure of the CDF $F_{y|\mathbf{x}}(\cdot)$.

Also, the ridge subspace is not unique for a particular $f(\mathbf{x})$. Consider $y = f(\mathbf{x}) = g(\boldsymbol{A}^T\mathbf{x})$ for some $\boldsymbol{A} \in \mathbb{R}^{m\times n}$ with $n < m$. Define $\boldsymbol{B} = \begin{bmatrix}\boldsymbol{A} & \mathbf{v}\end{bmatrix}$ with $\mathbf{v} \notin \text{colspan}(\boldsymbol{A})$. We can then write $y = f(\mathbf{x}) = \tilde{g}(\boldsymbol{B}^T\mathbf{x})$ where $\tilde{g}$ (i) maps the first n components of $\boldsymbol{B}^T\mathbf{x}$ as g and (ii) ignores the $(n+1)$th component.

However, since the goal is dimension reduction, we prefer to choose the ridge profile with the fewest inputs.

The minimum DRS from Definition 2 is the DRS which maximizes dimension reduction for a particular regression problem. For deterministic functions, we have the following definition of the *minimum ridge subspace*.

Definition 6 (Minimum ridge subspace)**.** *A subspace $\mathcal{S}$ is a* minimum ridge subspace *if* $\dim(\mathcal{S}) \leq \dim(\mathcal{S}_{\mathrm{R}})$ *for all ridge subspaces* $\mathcal{S}_{\mathrm{R}}$*.*

A minimum ridge subspace always exists due to the finite dimensionality of the input space. Thus, the notion of structural dimension from SDR extends to ridge functions. In this case, the structural dimension represents the fewest number of directions required to capture all of the variability in $f(\mathbf{x})$.

In section 2.1, we noted that the minimum DRS is not necessarily unique, and we defined the central DRS to obtain uniqueness. The central DRS is the intersection of all other DRSs, and it is the unique minimum DRS—provided it exists. For deterministic functions, we can define the *central ridge subspace* similarly.

Definition 7 (Central ridge subspace)**.** *A subspace $\mathcal{S}_{f(\mathbf{x})}$ is a* central ridge subspace *for $y = f(\mathbf{x})$ if $\mathcal{S}_{f(\mathbf{x})} \subseteq \mathcal{S}_{\mathrm{R}}$ for all ridge subspaces $\mathcal{S}_{\mathrm{R}}$.*

In words, when the central ridge subspace exists, it is the intersection of all ridge subspaces.

Theorems 1, 2, and 3 provide a variety of existence conditions for the central DRS in regression, but these conditions may be difficult or impossible to verify. In contrast, the deterministic function is accompanied by a known density function $\rho(\mathbf{x})$ on the inputs. We use the given $\rho(\mathbf{x})$ in an analog of Theorem 3 to derive existence conditions for the central ridge subspace; the proof is in section Appendix A.4. Assume that the support of the input density $\rho(\mathbf{x})$ is convex. Let $y = f(\mathbf{x})$, and assume f is differentiable for all $\mathbf{x}$ in the support of $\rho(\mathbf{x})$. Suppose that $\mathcal{S}_1$ and $\mathcal{S}_2$ are two ridge subspaces for $f(\mathbf{x})$. Then, $\mathcal{S}_1 \cap \mathcal{S}_2$ is also a ridge subspace.

Next, we extend Properties 1 and 2 for the central DRS to the central ridge subspace. These properties are essential to the development of the DRS approximation methods SIR and SAVE. To apply SIR and SAVE algorithms to dimension reduction in deterministic functions, we derive analogous properties for the central ridge subspace.

Property 1 relates the central DRSs for two regressions when their predictors are related by an affine transformation. A similar property holds for the central ridge subspace.

Property 3. *Let* $\mathbf{z} = \boldsymbol{B}\mathbf{x} + \mathbf{b}$ *be an affine transformation of inputs* $\mathbf{x}$*, where* $\boldsymbol{B} \in \mathbb{R}^{m \times m}$ *is full rank and* $\mathbf{b} \in \mathbb{R}^m$*. Given* $y = f(\mathbf{x})$*, define*

$$\tilde{f}(\mathbf{z}) \;=\; f\left(\boldsymbol{B}^{-1}(\mathbf{z} - \mathbf{b})\right). \tag{27}$$

Let $\mathcal{S}_{f(\mathbf{x})}$ *denote the central ridge subspace for* $f(\mathbf{x})$ *with basis matrix* $\boldsymbol{A}$*. Then,* $\text{colspan}(\boldsymbol{B}^{-T}\boldsymbol{A})$ *is the central ridge subspace for* $\tilde{f}(\mathbf{z})$*.*

A proof of Property 3 is in section Appendix A.5. This property simplifies analysis by allowing us to assume standardized inputs.

Assumption 2 (Standardized inputs)**.** *Assume that the inputs* $\mathbf{x}$ *in Problem 2 are standardized so that*

$$\int \mathbf{x}\, d\rho(\mathbf{x}) = \mathbf{0}, \qquad \int \mathbf{x}\,\mathbf{x}^T\, d\rho(\mathbf{x}) = \boldsymbol{I}. \tag{28}$$

For the remainder we assume that Assumption 2 holds unless otherwise stated.

Property 2 showed that the central DRS for $h(y)|\mathbf{x}$ is contained within the central DRS for $y|\mathbf{x}$ for any mapping of the response $h(y)$. A similar result holds in the deterministic case. A proof for this property is in section Appendix A.6.

Property 4. *Let* $\mathcal{S}_{f(\mathbf{x})}$ *be the central ridge subspace for* $y = f(\mathbf{x})$ *and* $h(y)$ *be any function of the output. Then,*

$$\mathcal{S}_{h(f(\mathbf{x}))} \subseteq \mathcal{S}_{f(\mathbf{x})}. \tag{29}$$

In words, the central ridge subspace for $h(y) = h(f(\mathbf{x}))$ is contained within the central ridge subspace for $y = f(\mathbf{x})$.

Next, we formalize *ridge recovery* as the search for the central ridge subspace for a given function and input density.

Problem 2 (Ridge recovery)**.** *Given an input probability density function* $\rho : \mathbb{R}^m \to \mathbb{R}^+$*, assume* $f : \mathbb{R}^m \to \mathbb{R}$ *is a ridge function of the form*

$$f(\mathbf{x}) \;=\; g(\boldsymbol{A}^T\mathbf{x}) \tag{30}$$

for some $\boldsymbol{A} \in \mathbb{R}^{m\times n}$ *with* $n < m$ *and* $g : \mathbb{R}^n \to \mathbb{R}$*. Using point evaluations of* $f(\mathbf{x})$ *for* $\mathbf{x}$ *in the support of* ρ*, compute a basis for the central ridge subspace* $\mathcal{S}_{f(\mathbf{x})} = \text{colspan}(\boldsymbol{A})$*.*

Recent papers in the signal processing and approximation theory literature propose and analyze algorithms for ridge recovery. Cohen et al. [27] consider the case of $n = 1$ with conditions on the ridge direction. Fornasier, Schnass, and Vybiral [17] expand on Cohen et al. and propose an algorithm for general $n < m$ under conditions on $\boldsymbol{A}$. Tyahi and Cevher [18] extend the approximation results from Fornasier et al. to more general $\boldsymbol{A}$'s. One goal of our paper is to interpret the inverse regression methods SIR and SAVE as methods for ridge recovery when applied to a given ridge function.

The ridge recovery problem is distinct from ridge approximation [28], where the goal is to find $\boldsymbol{A}$ and construct g that minimize the approximation error for a given $f(\mathbf{x})$. Inverse regression may be useful for ridge approximation or identifying ridge-like structure in a given function, but pursuing these ideas is beyond the scope of this manuscript. Li, Lin, and Li [15] distinguish between an "ideal scenario"—which loosely corresponds to ridge recovery and where they can prove certain convergence results for inverse regression—and the "general setting"—where they apply inverse regression complemented by heuristics to numerical examples.

It is worth discussing the role of the input probability density $\rho(\mathbf{x})$ in the ridge recovery problem. Lemma 3 places restrictions on the support of ρ to

ensure the existence of the central ridge subspace. And two different ρ's may produce different ridge subspaces—depending on f. We emphasize that the input density must is given in the ridge recovery problem; we cannot alter ρ in searching for the central ridge subspace. But the given ρ does impact the seach. In this paper, we use ρ to inform Monte Carlo sampling of $\mathbf{x}$ as we investigate direct applications of the SIR and SAVE algorithms to Problem 2.

Extending SDR theory to ridge functions provides a theoretical framework for studying how SIR (Algorithm 1) and SAVE (Algorithm 2) may be used for ridge recovery. Recall that for the regression, SIR and SAVE construct estimates for at least part of the central DRS using the inverse regression $\mathbf{x}|y$. The inverse regression of $y = f(\mathbf{x})$,

$$f^{-1}(y) = \{\, \mathbf{x} \in \mathbb{R}^m \,|\, y = f(\mathbf{x}) \,\}, \tag{31}$$

is the deterministic analog of the inverse regression $\mathbf{x}|y$. The inverse image is a subset of $\mathbb{R}^m$ parameterized by the output y. This differs from the inverse regression which is an m-dimensional conditional random vector parameterized by the response. For the inverse image to be useful in discovering the central ridge subspace, we must understand the connection between $\mathcal{S}_{f(\mathbf{x})}$ and any subspace we can construct using $f^{-1}(y)$.

3.1. SIR for ridge recovery

Recall from section 2.2 that the subspace containing inverse regression function $\mathbb{E}\,[\mathbf{x}|y]$ is contained in the central DRS; this fact motivates Algorithm 1. In this section, we study to what extent SIR may be used for ridge recovery (Problem 2).

We first consider the conditional expectation of $\mathbf{x}$ over the inverse image,

$$\boldsymbol{\mu}(y) \;=\; \int \mathbf{x}\, d\mu_{\mathbf{x}|y}(\mathbf{x}), \tag{32}$$

where $\mu_{\mathbf{x}|y}$ is the conditional probability measure on the set $f^{-1}(y)$ induced by the given probability density $\rho(\mathbf{x})$. Equation (32) is the average of the input $\mathbf{x}$

over the set $f^{-1}(y)$ with respect to ρ. If f is a ridge function, $f(\mathbf{x}) = g(\boldsymbol{A}^T\mathbf{x})$, then we may rewrite (32) as

$$\boldsymbol{\mu}(y) \;=\; \int \left[\int \mathbf{x}\, d\mu_{\mathbf{x}|\mathbf{u}}(\mathbf{x}) \right] d\mu_{\mathbf{u}|y}(\mathbf{u}) \tag{33}$$

where the conditional measures $\mu_{\mathbf{x}|\mathbf{u}}$ and $\mu_{\mathbf{u}|y}$ on the sets $\mathcal{X}(\mathbf{u}) = \left\{ \mathbf{x} \in \mathbb{R}^m \,|\, \mathbf{u} = \boldsymbol{A}^T\mathbf{x} \right\}$ and $g^{-1}(y) \;=\; \{\, \mathbf{u} \in \mathbb{R}^n \,|\, y = g(\mathbf{u}) \,\}$, respectively, are both appropriately constructed from the given density $\rho(\mathbf{x})$. Equation (33) decomposes the integral in the conditional expectation into two integrals. The outer integral averages over all $\mathbf{u} \in \mathbb{R}^n$ that map to a given y via g. The inner integral averages over all $\mathbf{x} \in \mathbb{R}^m$ that map to $\mathbf{u} \in g^{-1}(y)$. The decomposition of (32) into (33) is used in the analysis of SIR for ridge recovery.

In SIR for regression, the covariance of the inverse regression function is related to the central DRS. Recalling the standardization of $\mathbf{x}$ from Assumption 2, the analog of the covariance matrix for ridge recovery is

$$\boldsymbol{C}_{\text{SIR}} \;=\; \int \boldsymbol{\mu}(y)\, \boldsymbol{\mu}(y)^T\, d\gamma(y) \tag{34}$$

where γ is the measure on $\mathbb{R}$ induced by $\rho(\mathbf{x})$ and $y = f(\mathbf{x})$. Equation (34) is the average over y of the outer product of the conditional average $\boldsymbol{\mu}(y)$ with itself. The following lemma provides insight into the null space of $\boldsymbol{C}_{\text{SIR}}$; the proof is in section Appendix A.7. A vector $\mathbf{w} \in \mathbb{R}^m$ is in the null space of $\boldsymbol{C}_{\text{SIR}}$ if and only if $\boldsymbol{\mu}(y)^T\mathbf{w} = 0$ for all y in the range of f. Lemma 3.1 is the key to the proof of the next theorem that connects $\text{colspan}(\boldsymbol{C}_{\text{SIR}})$ to the central ridge subspace; the proof is in section Appendix A.8.

Theorem 7. *Given an elliptically symmetric input probability density function $\rho : \mathbb{R}^m \to \mathbb{R}^+$, assume $f : \mathbb{R}^m \to \mathbb{R}$ admits a central ridge subspace $\mathcal{S}_{f(\mathbf{x})}$. Then,* $\text{colspan}(\boldsymbol{C}_{\text{SIR}}) \subseteq \mathcal{S}_{f(\mathbf{x})}$.

Theorem 7 implies that $\boldsymbol{C}_{\text{SIR}}$ can be used to uncover at least part of f's central ridge subspace $\mathcal{S}_{f(\mathbf{x})}$. In the following example, $\text{colspan}(\boldsymbol{C}_{\text{SIR}})$ is a strict subset of $\mathcal{S}_{f(\mathbf{x})}$.

Example 3. *Assume* $\mathbf{x} \in \mathbb{R}^2$*, and let* $\rho(\mathbf{x})$ *be a bivariate standard Gaussian. Let* $y = f(\mathbf{x}) = x_1 x_2$*. For any value of* y*,* $\mathbf{x} \in f^{-1}(y)$ *implies* $-\mathbf{x} \in f^{-1}(y)$*. Therefore,* $\boldsymbol{\mu}(y) = \mathbf{0}$ *for all* y*, and* $\boldsymbol{C}_{\text{SIR}} = \mathbf{0}$*. But* y *is not constant over* $\mathbb{R}^2$*. Thus,* $\{\mathbf{0}\} = \text{colspan}(\boldsymbol{C}_{\text{SIR}}) \subset \mathcal{S}_{f(\mathbf{x})} = \mathbb{R}^2$*.*

Example 3 shows that there are functions that are not ridge functions for which $\boldsymbol{C}_{\text{SIR}}$ is low-rank. Therefore, discovering rank deficiency in $\boldsymbol{C}_{\text{SIR}}$ for a given f is not sufficient to identify f as a ridge function.

Theorem 7 requires that the input density be elliptically symmetric, which is satisfied when $\rho(\mathbf{x})$ is a multivariate normal. However, uncertain inputs in many engineering models are not normally distributed. Cook and Nachtsheim [29] develop a method for reweighting a set of samples such that they appear to be drawn from an elliptically symmetric distribution. This reweighting can improve the performance of SIR, and it does not alter the central ridge subspace.

Property 4 shows that the central ridge subspace for any mapping of y is contained within the central ridge subspace for y. That is, $\mathcal{S}_{h(f(\mathbf{x}))} \subseteq \mathcal{S}_{f(\mathbf{x})}$ for any $h(y)$. If we let $h(y)$ be the slicing function from (14), then the continuous integral form of $\boldsymbol{C}_{\text{SIR}}$ becomes a summation,

$$\boldsymbol{D}_{\text{SIR}} \equiv \sum_{r=1}^{R} \omega(r)\, \boldsymbol{\mu}(r)\, \boldsymbol{\mu}(r)^T, \tag{35}$$

where

$$\omega(r) = \int d\mu_{\mathbf{x}|r}(\mathbf{x}) = \mathbb{P}[f(\mathbf{x}) \in J_r] \tag{36}$$

is a probability mass function on $\{1, \ldots, R\}$; the $\mu_{\mathbf{x}|r}$ is the conditional measure induced by $\rho(\mathbf{x})$ and the composite map $h \circ f$. The slice expectation $\boldsymbol{\mu}(r)$ is the average of $\mathbf{x}$ over all $\mathbf{x}$ for which $f(\mathbf{x}) \in J_r$,

$$\boldsymbol{\mu}(r) = \int \mathbf{x}\, d\mu_{\mathbf{x}|r}(\mathbf{x}). \tag{37}$$

Applying Property 4 to the sliced mapping yields

$$\text{colspan}(\boldsymbol{D}_{\text{SIR}}) \subseteq \mathcal{S}_{h(f(\mathbf{x}))} \subseteq \mathcal{S}_{f(\mathbf{x})}. \tag{38}$$

Algorithm 1 uses N point evaluations of f to construct $\hat{\boldsymbol{D}}_{\text{SIR}}$ that approximates $\boldsymbol{D}_{\text{SIR}}$ from (35). Therefore, the SIR algorithm can be used to estimate a solution to Problem 2.

3.1.1. Convergence of SIR

Next, we study the convergence properties of the SIR algorithm for ridge recovery. Recall from Algorithm 1 that N_r denotes the number of samples in each slice J_r for $r = 1, \ldots, R$. The convergence depends on the smallest number of samples per slice over all the slices. We denote this value by

$$N_{r_{\min}} \;=\; \min_{1 \le r \le R} N_r. \tag{39}$$

Also, recall that the slices are assumed to be constructed such that $\omega(r) = \mathbb{P}[y \in J_r] > 0$. Thus, $N_{r_{\min}} > 0$ with probability 1 as $N \to \infty$. Note that the matrix $\hat{\boldsymbol{D}}_{\text{SIR}}$ is a random matrix, since the $\mathbf{x}_i$ are drawn randomly according to $\rho(\mathbf{x})$. Therefore, anything computed using $\hat{\boldsymbol{D}}_{\text{SIR}}$—including eigenvalues and eigenvectors—is also random. Therefore, the convergence analysis is probabilistic.

The following theorem shows that the eigenvalues of $\hat{\boldsymbol{D}}_{\text{SIR}}$ converge to those of $\boldsymbol{D}_{\text{SIR}}$ in a mean-squared sense; the proof can be found in section Appendix A.9.

Theorem 8. *Assume that Algorithm 1 has been applied to the data set* $\{[\, y_i \,,\, \mathbf{x}_i^T \,]\}$, *with* $i = 1, \ldots, N$, *where the* $\mathbf{x}_i$ *are drawn independently according to* $\rho(\mathbf{x})$ *and* $y_i = f(\mathbf{x}_i)$ *are point evaluations of* $f(\mathbf{x})$. *Then, for* $k = 1, \ldots, m$,

$$\mathbb{E}\left[\left(\lambda_k(\boldsymbol{D}_{\text{SIR}}) - \lambda_k(\hat{\boldsymbol{D}}_{\text{SIR}})\right)^2\right] \;=\; \mathcal{O}(N_{r_{\min}}^{-1}) \tag{40}$$

where $\lambda_k(\cdot)$ *denotes the* k*th eigenvalue of the given matrix.*

In words, the mean-squared error in the eigenvalues of $\hat{\boldsymbol{D}}_{\text{SIR}}$ decays at a $N_{r_{\min}}^{-1}$ rate. Since $\omega(r) > 0$ for all r, $N_{r_{\min}} \to \infty$ as $N \to \infty$. Moreover, the convergence rate suggests that one should define the slices in Algorithm 1 such that the same number of samples appears in each slice. This maximizes $N_{r_{\min}}$ and reduces the error in the eigenvalues.

The next theorem shows the convergence of the subspaces produced by Algorithm 1 as the number of point evaluations of f increases. We measure convergence using the subspace distance [30],

$$\text{dist}\left(\text{colspan}(\boldsymbol{A}), \text{colspan}(\hat{\boldsymbol{A}})\right) \;=\; \left\|\boldsymbol{A}\boldsymbol{A}^T - \hat{\boldsymbol{A}}\hat{\boldsymbol{A}}^T\right\|_2, \tag{41}$$

where $\boldsymbol{A}, \hat{\boldsymbol{A}}$ are the first n eigenvectors of $\boldsymbol{D}_{\text{SIR}}$ and $\hat{\boldsymbol{D}}_{\text{SIR}}$, respectively. The distance metric (41) is the principal angle between the subspaces colspan($\boldsymbol{A}$) and colspan($\hat{\boldsymbol{A}}$). The proof for the following theorem is in section Appendix A.10.

Theorem 9. *Assume the same conditions from Theorem 8. Then, for sufficiently large N,*

$$\text{dist}\left(\text{colspan}(\boldsymbol{A}), \text{colspan}(\hat{\boldsymbol{A}})\right) \;=\; \frac{1}{\lambda_n(\boldsymbol{D}_{\text{SIR}}) - \lambda_{n+1}(\boldsymbol{D}_{\text{SIR}})}\,\mathcal{O}(N_{r_{\min}}^{-1}) \tag{42}$$

with high probability.

The subspace error decays at a rate $N_{r_{\min}}^{-1}$ with high probability for sufficiently large N. The error in the n-dimensional approximate subspace depends inversely on the gap between the nth and $(n+1)$th eigenvalues of $\boldsymbol{D}_{\text{SIR}}$. In practice, the eigenvalues of $\hat{\boldsymbol{D}}_{\text{SIR}}$ provide insight into the subspace error: a large gap between eigenvalues suggests a better estimate of the subspace for a fixed number of samples.

3.2. SAVE for ridge recovery

In section 2.3, we discussed how SAVE for regression uses the covariance matrix of the inverse regression to estimate at least part of the central DRS. Algorithm 2 constructs a sample approximation of the SAVE subspace. Here, we study how the SAVE algorithm can be used to discover the central ridge subspace.

The analog of the conditional covariance matrix of $\mathbf{x}$ is

$$\boldsymbol{\Sigma}(y) \;=\; \int (\mathbf{x} - \boldsymbol{\mu}(y))\,(\mathbf{x} - \boldsymbol{\mu}(y))^T \; d\mu_{\mathbf{x}|y}(\mathbf{x}), \tag{43}$$

where $\boldsymbol{\mu}(y)$ is from (33). Equation (43) is a measure of the spread of the inverse image. This matrix may be rewritten in the form

$$\boldsymbol{\Sigma}(y) \;=\; \int \mathbf{x}\mathbf{x}^T \, d\mu_{\mathbf{x}|y}(\mathbf{x}) - \boldsymbol{\mu}(y)\boldsymbol{\mu}(y)^T, \tag{44}$$

which is useful when analyzing the SAVE algorithm for ridge recovery.

The essential matrix in SAVE for regression is $\mathbb{E}\left[(\boldsymbol{I} - \mathrm{Cov}\left[\mathbf{x}|y\right])^2\right]$; see section 2.3. The analog for ridge recovery is

$$\boldsymbol{C}_{\mathrm{SAVE}} \;=\; \int \left[\boldsymbol{I} - \boldsymbol{\Sigma}(y)\right]^2 \, d\gamma(y), \tag{45}$$

where γ is as in (34). The following lemma provides an interpretation of the null space of $\boldsymbol{C}_{\mathrm{SAVE}}$; the proof for this lemma is in section Appendix A.11. A vector $\mathbf{w}$ is in the null space of $\boldsymbol{C}_{\mathrm{SAVE}}$ if and only if $\boldsymbol{\Sigma}(y)\mathbf{w} = \mathbf{w}$ for all y in the range of f. Lemma 3.2 is used in the proof of the following theorem which connects $\mathrm{colspan}(\boldsymbol{C}_{\mathrm{SAVE}})$ to the central ridge subspace; the proof is in section Appendix A.12.

Theorem 10. *Given an elliptically symmetric input probability density function $\rho : \mathbb{R}^m \to \mathbb{R}^+$, assume $f : \mathbb{R}^m \to \mathbb{R}$ admits a central ridge subspace $\mathcal{S}_{f(\mathbf{x})}$. Then,* $\mathrm{colspan}(\boldsymbol{C}_{\mathrm{SAVE}}) \subseteq \mathcal{S}_{f(\mathbf{x})}$.

Theorem 10 implies that the column space of $\boldsymbol{C}_{\mathrm{SAVE}}$ can be used to uncover at least part of the central ridge subspace. The following example shows that these subspaces are not equal, in general.

Example 4. *Assume $\mathbf{x} \in \mathbb{R}^2$, and let $\rho(\mathbf{x})$ be a bivariate standard Gaussian. Let $y = f(\mathbf{x}) = \frac{1}{\sqrt{2}}x_1^2 + x_2$. Note,*

$$\int \mathbf{x}\mathbf{x}^T \, d\mu_{\mathbf{x}|y}(\mathbf{x}) \;=\; \begin{bmatrix} \sqrt{2}y & 0 \\ 0 & y^2 - \sqrt{2}y + \frac{1}{2} \end{bmatrix} \tag{46}$$

and

$$\boldsymbol{\mu}(y) \;=\; \int \mathbf{x} \, d\mu_{\mathbf{x}|y}(\mathbf{x}) \;=\; \begin{bmatrix} 0 \\ y - \frac{1}{\sqrt{2}} \end{bmatrix}. \tag{47}$$

Then

$$\begin{aligned}\boldsymbol{\Sigma}(y) &= \int \mathbf{x}\mathbf{x}^T \, d\mu_{\mathbf{x}|y}(\mathbf{x}) - \boldsymbol{\mu}(y)\,\boldsymbol{\mu}(y)^T \\ &= \begin{bmatrix} \sqrt{2}y & 0 \\ 0 & y^2 - \sqrt{2}y + \frac{1}{2} \end{bmatrix} - \begin{bmatrix} 0 \\ y - \frac{1}{\sqrt{2}} \end{bmatrix} \begin{bmatrix} 0 & y - \frac{1}{\sqrt{2}} \end{bmatrix} \\ &= \begin{bmatrix} \sqrt{2}y & 0 \\ 0 & 1 \end{bmatrix}.\end{aligned} \tag{48}$$

Let $\mathbf{b} = [\,0\,,\,1\,]^T$ *and note*

$$\boldsymbol{\Sigma}(y)\,\mathbf{b} = \begin{bmatrix} \sqrt{2}y & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \end{bmatrix} = \mathbf{b}, \tag{49}$$

for all y. *By Lemma 3.2,* $\mathbf{b} \in \text{null}(\boldsymbol{C}_{\text{SAVE}})$ *which implies that* $\boldsymbol{C}_{\text{SAVE}}$ *is rank-1. However, the central ridge subspace of* $f(\mathbf{x})$ *is all of* $\mathbb{R}^2$. *Hence,* $\text{colspan}(\boldsymbol{C}_{\text{SAVE}})$ *is a strict subset of* $\mathcal{S}_{f(\mathbf{x})}$.

Example 4 shows that there are functions that are not ridge functions for which $\boldsymbol{C}_{\text{SAVE}}$ is low-rank. Therefore, discovering rank deficiency in $\boldsymbol{C}_{\text{SAVE}}$ for a given f is not sufficient to identify f as a ridge function.

By Property 4, for the slicing map h from (14), the continuous integral form of $\boldsymbol{C}_{\text{SAVE}}$ becomes a summation,

$$\boldsymbol{D}_{\text{SAVE}} \;\equiv\; \sum_{r=1}^{R} \omega(r)\,[\boldsymbol{I} - \boldsymbol{\Sigma}(r)]^2\,, \tag{50}$$

where $\omega(r)$ is from (36). The slice covariance $\boldsymbol{\Sigma}(r)$ is a measure of the spread of inputs samples that map via the composite function $h \circ f$ to the rth slice,

$$\boldsymbol{\Sigma}(r) \;=\; \int (\mathbf{x} - \boldsymbol{\mu}(r))\,(\mathbf{x} - \boldsymbol{\mu}(r))^T \; d\mu_{\mathbf{x}|r}(\mathbf{x}) \tag{51}$$

where $\boldsymbol{\mu}(r)$ is from (37). By Property 4, $\text{colspan}(\boldsymbol{D}_{\text{SAVE}})$ is contained in the central ridge subspace,

$$\text{colspan}(\boldsymbol{D}_{\text{SAVE}}) \;\subseteq\; \mathcal{S}_{h(f(\mathbf{x}))} \;\subseteq\; \mathcal{S}_{f(\mathbf{x})}. \tag{52}$$

Algorithm 2 approximates $\boldsymbol{D}_{\text{SAVE}}$ using the data set of point evaluations $\{[\,y_i\,,\,\mathbf{x}_i^T]\}$, $i = 1, \ldots, N$. This implies that the SAVE algorithm can recover at least part of the central ridge subspace.

3.2.1. Convergence of SAVE

Next, we study the convergence of SAVE for ridge recovery. Recall $N_{r_{\min}}$ from (39). The following theorem shows the rate of mean-squared convergence of the eigenvalues of $\hat{\boldsymbol{D}}_{\text{SAVE}}$; the proof is in section Appendix A.13.

Theorem 11. *Assume that Algorithm 2 has been applied to the data set* $\{[\, y_i \,,\, \mathbf{x}_i^T \,]\}$*, with* $i = 1, \ldots, N$*, where the* $\mathbf{x}_i$ *are drawn independently according to* $\rho(\mathbf{x})$ *and* $y_i = f(\mathbf{x}_i)$ *are point evaluations of* $f(\mathbf{x})$*. Then, for* $k = 1, \ldots, m$*,*

$$\mathbb{E}\left[\left(\lambda_k(\boldsymbol{D}_{\text{SAVE}}) - \lambda_k(\hat{\boldsymbol{D}}_{\text{SAVE}})\right)^2\right] \;=\; \mathcal{O}(N_{r_{\min}}^{-1}) \tag{53}$$

where $\lambda_k(\cdot)$ *denotes the* k*th eigenvalue of the given matrix.*

In words, the mean-squared error in the eigenvalues of $\hat{\boldsymbol{D}}_{\text{SAVE}}$ decays at a rate of $N_{r_{\min}}^{-1}$. This highlights the value of defining the slices such that they maximize $N_{r_{\min}}$ in a given data set.

Next, we examine the convergence of the subspaces Algorithm 2 produces, where the subspace distance is from (41). A proof for the following theorem can be found in section Appendix A.14.

Theorem 12. *Assume the same conditions from Theorem 11. Then, for sufficiently large* N*,*

$$\text{dist}\left(\text{colspan}(\boldsymbol{A}),\, \text{colspan}(\hat{\boldsymbol{A}})\right) \;=\; \frac{1}{\lambda_n(\boldsymbol{D}_{\text{SAVE}}) - \lambda_{n+1}(\boldsymbol{D}_{\text{SAVE}})}\, \mathcal{O}(N_{r_{\min}}^{-1}) \tag{54}$$

with high probability.

The subspace error for Algorithm 2 decays like $N_{r_{\min}}^{-1}$ with high probability for sufficiently large N. Similar to the estimated SIR subspace from Algorithm 1, the error depends inversely on the eigenvalue gap. If the gap between the nth and $(n+1)$th eigenvalues is large, then the error in the estimated n-dimensional subspace is small for a fixed number of samples.

3.3. Interpreting the eigenvalues

In sections 3.1 and 3.2, we explained how the SIR and SAVE algorithms attempt to approximate the central ridge subspace. In this section, we highlight

how this approximation occurs indirectly through the SIR and SAVE subspaces. These subspaces (i) are properties of the given function and (ii) are subspaces of the central ridge subspace—see (38) and (52). Algorithms 1 and 2 approximate the SIR and SAVE subspaces, respectively, to gain insight into the central ridge subspace. The convergence analysis in sections 3.1.1 and 3.2.1 shows that increasing the number of samples improves the approximation of the SIR and SAVE subspace; however, this does not address how accurately the SIR and SAVE subspaces approximate the central ridge subspace.

The subspace containments from (38) and (52) also provide insight into the eigenvalues of the SIR and SAVE matrices. Eigenvectors associated with nonzero eigenvalues define the n-dimensional estimated SIR and SAVE subspaces; however, these eigenvalues may incorrectly suggest ridge structure that is not present. Examples 3 and 4 show functions that produce such false positives for each algorithm. These are cases when the SIR or SAVE subspace being approximated does not equal the central ridge subspace.

Eigenvalues from Algorithms 1 and 2 estimate the true eigenvalues of the SIR and SAVE matrices, respectively. For a finite number of samples, the eigenvalue estimates almost never equal zero even if the true eigenvalue does. This raises the issue of determining the dimensionality of the estimated subspace. This issue closely relates to the idea of discovering ridge-like structure in a function. Both of these topics are outside the scope of this paper as we focus on these methods' abilities to solve the ridge recovery problem. For this reason, we treat n—the dimension of the ridge subspace—as given in the SIR and SAVE algorithms.

4. Numerical results

We examine the performance of the SIR and SAVE algorithms for ridge recovery in three problems that exhibit known ridge structure. The first two problems are multivariate quadratic functions, and the third problem is a physically-motivated test problem from magnetohydrodynamics. We provide several graphics in the numerical study of these problems including plots of estimated eigen-

values, eigenvalue errors, and subspace errors. One graphic is especially useful for visualizing the structure of the function relative to its central ridge subspace: the sufficient summary plot [8]. Sufficient summary plots show y versus $\boldsymbol{A}^T\mathbf{x}$, where $\boldsymbol{A}$ contains a basis for the central ridge subspace. Algorithms 1 and 2 produce eigenvectors that span an approximation of the SIR and SAVE subspaces, respectively. We use these eigenvectors to construct the low-dimensional inputs for sufficient summary plots.

The python code used to generate the figures found throughout this section are available at `https://bitbucket.org/aglaws/inverse-regression-for-ridge-recovery`. The scripts require the dev branch of the Python Active-subspaces Utility Library, which can be found at `https://github.com/paulcon/active_subspaces/tree/dev`. We note that the convergence analysis from sections 3.1.1 and 3.2.1 assume a fixed slicing of the observed y range. However, Algorithms 1 and 2 are implemented in the utility library using an adaptive slicing approach that attempts to maximize $N_{r_{\min}}$ for a given set of data. This is done as a heuristic technique for reducing eigenvalue and subspace errors.

4.1. One-dimensional quadratic function

Let $\mathbf{x} \in \mathbb{R}^{10}$, and let $\rho(\mathbf{x})$ be a standard normal density function. Define the function

$$y \;=\; f(\mathbf{x}) \;=\; \left(\mathbf{b}^T\mathbf{x}\right)^2, \tag{55}$$

where $\mathbf{b} \in \mathbb{R}^{10}$ is a constant vector. The span of $\mathbf{b}$ is the central ridge subspace. First, we attempt to uncover the central ridge subspace using SIR (Algorithm 1), which is known to fail for functions symmetric about $\mathbf{x} = \mathbf{0}$ [11]; Figure 1 confirms this failure. In fact, $\boldsymbol{C}_{\text{SIR}}$ for this problem is zero since the conditional expectation of $\mathbf{x}$ for any value of y is zero. Figure 1a shows that all estimated eigenvalues of the SIR matrix are nearly zero as expected. Figure 1b is a one-dimensional sufficient summary plot of y against $\hat{\mathbf{w}}_1^T\mathbf{x}$, where $\hat{\mathbf{w}}_1$ denotes the normalized eigenvector associated with the largest eigenvalue of $\hat{\boldsymbol{D}}_{\text{SIR}}$ from Algorithm 1. If (i) the central ridge subspace is one-dimensional (as in this case) and (ii) the chosen SDR algorithm correctly identifies the one basis vector, then

the sufficient summary plot will show a univariate relationship between the linear combination of input evaluations and the associated outputs. Due to the symmetry in the quadratic function, SIR fails to recover the basis vector. However, it should be noted that if (55) had the form $y = f(\mathbf{x}) = (\mathbf{b}^T\mathbf{x} + \mathbf{c})^2$ for a constant $\mathbf{c} \neq \mathbf{0}$, then SIR would not suffer these issues. This highlights that the function must be symmetric about zero for SIR to fail.

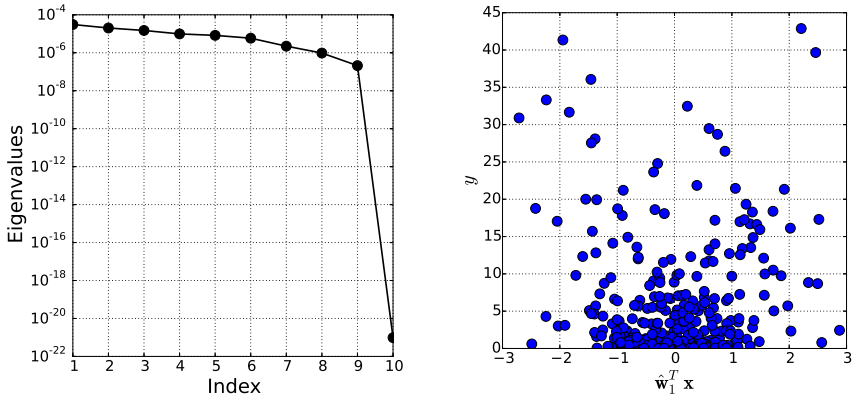

(a) Eigenvalues of $\hat{\boldsymbol{D}}_{\text{SIR}}$ from (11)

(b) Sufficient summary plot for SIR

Figure 1: As expected, SIR fails to recover the ridge direction $\mathbf{b}$ in the function (55).

Figure 2 shows results from applying SAVE (Algorithm 2) to the quadratic function (55). Figure 2a shows the eigenvalues of $\hat{\boldsymbol{D}}_{\text{SAVE}}$ from Algorithm 2. Note the large gap between the first and second eigenvalues, which suggests that the SAVE subspace is one-dimensional. Figure 2b shows the sufficient summary plot using the first eigenvector $\hat{\mathbf{w}}_1$ from Algorithm 2, which reveals the univariate quadratic relationship between $\hat{\mathbf{w}}_1^T\mathbf{x}$ and y.

4.2. Three-dimensional quadratic function

Next, we numerically study the convergence properties of the SIR and SAVE algorithms using a more complex quadratic function. Let $\mathbf{x} \in \mathbb{R}^{10}$, and let $\rho(\mathbf{x})$

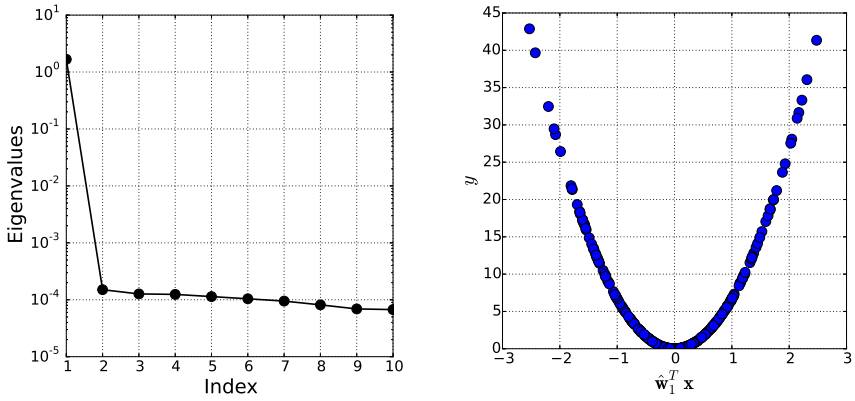

(a) Eigenvalues of $\hat{\boldsymbol{D}}_{\text{SAVE}}$ from (20)

(b) Sufficient summary plot for SAVE

Figure 2: SAVE recovers the ridge direction $\mathbf{b}$ in the function (55).

be a standard normal density. Define the function

$$y \;=\; f(\mathbf{x}) \;=\; \mathbf{x}^T\boldsymbol{B}\boldsymbol{B}^T\mathbf{x} + \mathbf{b}^T\mathbf{x}, \tag{56}$$

where $\boldsymbol{B} \in \mathbb{R}^{10\times 2}$ and $\mathbf{b} \in \mathbb{R}^{10}$ with $\mathbf{b} \notin \text{colspan}(\boldsymbol{B})$. We expect better results with SIR compared to the example in (55) since (56) is not symmetric about $\mathbf{x} = \mathbf{0}$. Figure 3a shows the eigenvalues decay of $\hat{\boldsymbol{D}}_{\text{SIR}}$; note the gap between the third and fourth eigenvalues. Figure 3b shows the maximum squared eigenvalue error normalized by the largest eigenvalue,

$$\max_{1\le i\le m} \frac{\Big(\lambda_i(\hat{\boldsymbol{D}}_{\text{SIR}}) - \lambda_i(\boldsymbol{D}_{\text{SIR}})\Big)^2}{\lambda_1(\boldsymbol{D}_{\text{SIR}})^2}, \tag{57}$$

for increasing numbers of samples in 10 independent trials. We estimate the true eigenvalues using SIR with 10^7 samples. The average error decays at a rate slightly faster than the $\mathcal{O}(N^{-1})$ from Theorem 8. The improvement can likely be attributed to the adaptive slicing procedure discussed at the beginning of this section. Figure 3c shows the error in the estimated three-dimensional SIR subspace (see (41)) for increasing numbers of samples in 10 independent trials. We use 10^7 samples to estimate the true SIR subspace. The subspace errors

decrease asymptotically at a rate of approximately $\mathcal{O}(N^{-1})$, which agrees with Theorem 9.

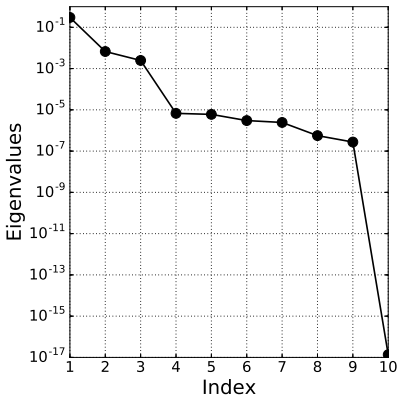

(a) Eigenvalues of $\hat{\boldsymbol{D}}_{\text{SIR}}$ from (11)

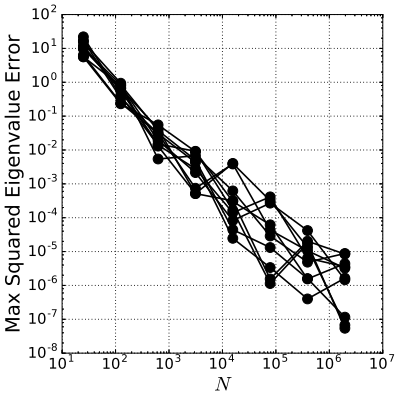

(b) Maximum squared eigenvalue error, SIR

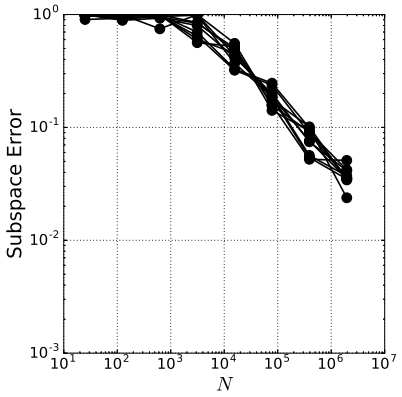

(c) SIR subspace errors for $n = 3$

Figure 3: Eigenvalues, eigenvalue errors, and subspace errors for SIR applied to (56). The error decreases with increasing samples consistent with the convergence theory in section 3.1.1.

Figure 4 shows the results of a similar convergence study using SAVE (Algorithm 2). The eigenvalues of $\hat{\boldsymbol{D}}_{\text{SAVE}}$ from (20) are shown in Figure 4a. Note the large gap between the third and fourth eigenvalues, which is consistent with the three-dimensional central ridge subspace in $f(\mathbf{x})$ from (56). Figures 4b and 4c

show the maximum squared eigenvalue error and the subspace error for $n = 3$, respectively. The eigenvalue error again decays at a faster rate than expected in Theorem 11 due to the adaptive slicing implemented in the code. The subspace error decays consistently according to Theorem 12.

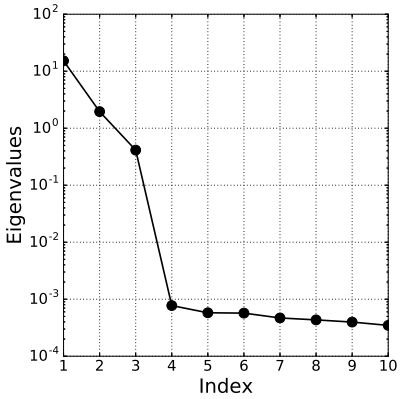

(a) Eigenvalues of $\hat{\boldsymbol{D}}_{\text{SAVE}}$ from (20)

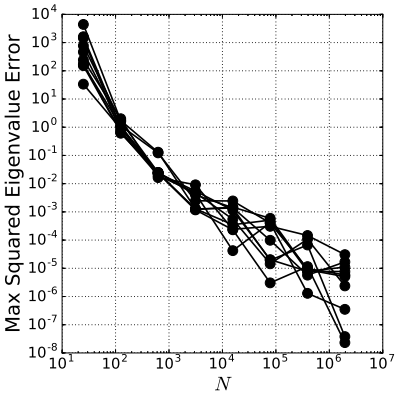

(b) Maximum squared eigenvalue error, SAVE

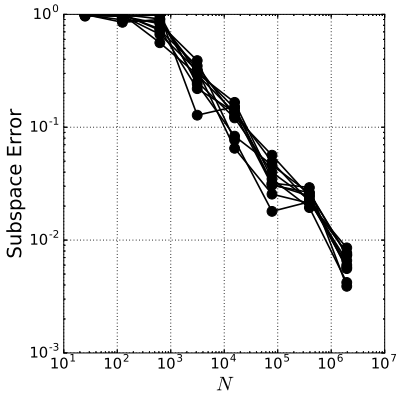

(c) SAVE subspace errors for $n = 3$

Figure 4: Eigenvalues, eigenvalue errors, and subspace errors for SAVE applied to (56). The error decreases with increasing samples consistent with the convergence theory in section 3.2.1.

4.3. Hartmann problem

We use the following model as a test case for parameter space dimension reduction methods in recent work; the problem set up here closely follows our previous development [31]. The Hartmann problem is a standard problem in magnetohydrodynamics (MHD) that models the flow of an electrically-charged plasma in the presence of a uniform magnetic field [32]. The flow occurs along an infinite channel between two parallel plates separated by distance 2ℓ. The applied magnetic field is perpendicular to the flow direction and acts as a resistive force on the flow velocity. At the same time, the movement of the fluid induces a magnetic field along the direction of the flow. Figure 5 contains a diagram of this problem.

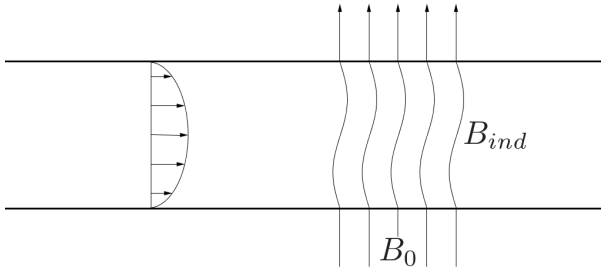

Figure 5: The Hartmann problem studies the flow of an ionized fluid between two parallel plates. A magnetic field applied perpendicular to the flow direction acts as a resistive force to the fluid flow. Simultaneously, the fluid induces a magnetic field along the direction of the flow.

The inputs to the Hartmann model are fluid viscosity μ, fluid density ρ, applied pressure gradient $\partial p_0/\partial x$ (where the derivative is with respect to the flow field's spatial coordinate), resistivity η, and applied magnetic field B_0. We collect these inputs into a vector,

$$\mathbf{x} = \begin{bmatrix} \mu & \rho & \frac{\partial p_0}{\partial x} & \eta & B_0 \end{bmatrix}^T. \tag{58}$$

We examine one output of interest: the total induced magnetic field

$$B_{\text{ind}}(\mathbf{x}) = \frac{\partial p_0}{\partial x} \frac{\ell \mu_0}{2B_0} \left(1 - 2\frac{\sqrt{\eta\mu}}{B_0 \ell} \tanh\left(\frac{B_0 \ell}{2\sqrt{\eta\mu}} \right) \right). \tag{59}$$

This function is not a ridge function of $\mathbf{x}$. However, it has been shown that many physical laws can be expressed as ridge functions by considering a log transform of the inputs [33]. For this reason, we apply SIR and SAVE ridge recovery to B_{ind} as a function of the logarithms of the inputs from (58). The log-transformed inputs are drawn from a multivariate Gaussian with

$$\boldsymbol{\mu} = \begin{bmatrix} -2.25 \\ 1 \\ 0.3 \\ 0.3 \\ -0.75 \end{bmatrix}, \qquad \boldsymbol{\Sigma} = \begin{bmatrix} 0.15 & & & & \\ & 0.25 & & & \\ & & 0.25 & & \\ & & & 0.25 & \\ & & & & 0.25 \end{bmatrix}. \tag{60}$$

Figure 6 shows the results of applying SIR (Algorithm 1) to the Hartmann model for the induced magnetic field B_{ind}. The eigenvalues of $\hat{\boldsymbol{D}}_{\text{SIR}}$ from (11) with bootstrap ranges are shown in Figure 6a. Large gaps appear after the first and second eigenvalues which indicate possible two-dimensional ridge structure. In fact, the induced magnetic field admits a two-dimensional central ridge subspace relative to the log-inputs [31]. Figures 6c and 6d contain one- and two-dimensional sufficient summary plots of B_{ind} against the $\hat{\mathbf{w}}_1^T\mathbf{x}$ and $\hat{\mathbf{w}}_2^T\mathbf{x}$, where $\hat{\mathbf{w}}_1$ and $\hat{\mathbf{w}}_2$ are the first two eigenvectors of $\hat{\boldsymbol{D}}_{\text{SIR}}$. We see a strong one-dimensional relationship. However, the structural dimension of B_{ind} is 2; the two-dimensional plot shows slight curvature with changes in $\hat{\mathbf{w}}_2^T\mathbf{x}$. These results suggest that ridge-like structure may be discovered using the SIR algorithm in some cases. Figure 6b shows the subspace errors as a function of the subspace dimension. Recall from Theorem 9 that the subspace error depends inversely on the eigenvalue gap. The largest eigenvalue gap occurs between the first and second eigenvalues, which is consistent with the smallest subspace error for $n = 1$.

We perform the same numerical studies for SAVE applied to B_{ind}. Figure 7a shows the eigenvalues of the $\hat{\boldsymbol{D}}_{\text{SAVE}}$ from (20) for the induced magnetic field B_{ind} from (59). Note the large gaps after the first and second eigenvalues. These gaps are consistent with the subspace errors in Figure 7b, where the one- and

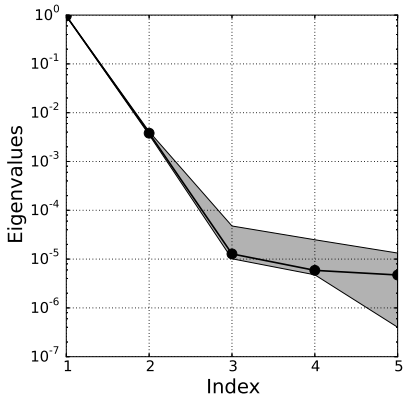

(a) Eigenvalues of $\hat{\boldsymbol{D}}_{\text{SIR}}$ with bootstrap ranges

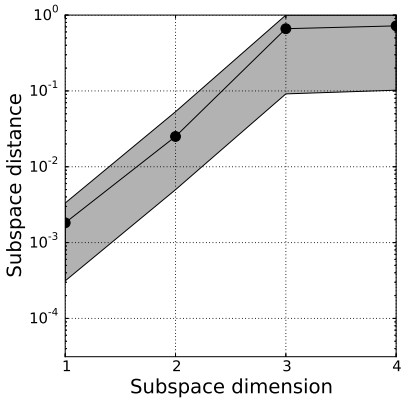

(b) Subspace errors from $\hat{\boldsymbol{D}}_{\text{SIR}}$

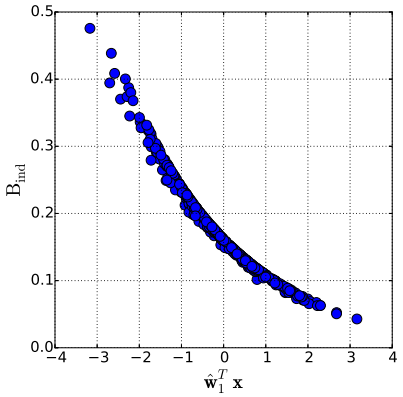

(c) One-dimensional SIR summary plot for B_{ind} from (59)

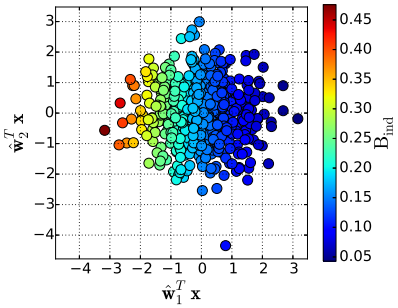

(d) Two-dimensional SIR summary plot for B_{ind} from (59)

Figure 6: Eigenvalues with bootstrap ranges, estimated subspace errors, and sufficient summary plots for SIR (Algorithm 1) applied to B_{ind} from (59).

two-dimensional subspace estimates have the smallest errors. Figures 7c and 7d contain sufficient summary plots for $\hat{\mathbf{w}}_1^T\mathbf{x}$ and $\hat{\mathbf{w}}_2^T\mathbf{x}$, where $\hat{\mathbf{w}}_1$ and $\hat{\mathbf{w}}_2$ are the first two eigenvectors from $\hat{\boldsymbol{D}}_{\text{SAVE}}$ in (20).

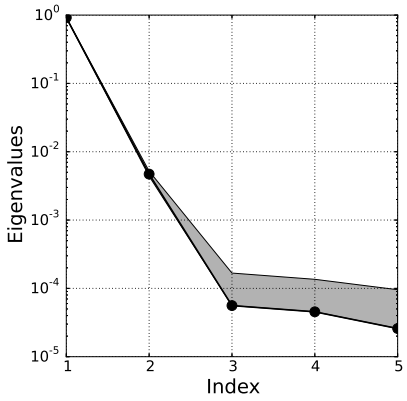

(a) Eigenvalues of $\hat{\boldsymbol{D}}_{\mathrm{SAVE}}$ with bootstrap ranges

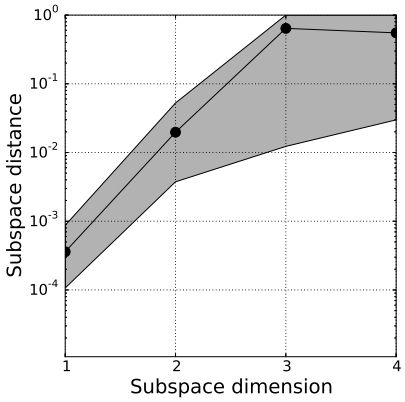

(b) Subspace errors from $\hat{\boldsymbol{D}}_{\mathrm{SAVE}}$

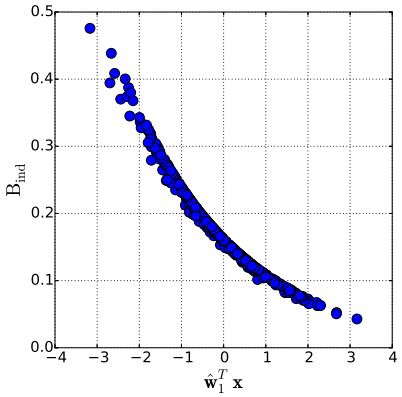

(c) One-dimensional SAVE summary plot for B_{ind} from (59)

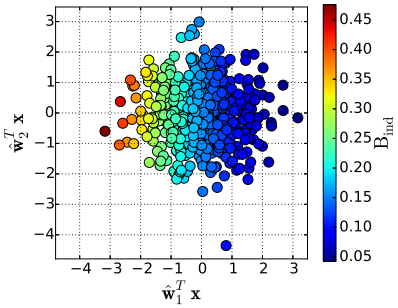

(d) Two-dimensional SAVE summary plot for B_{ind} from (59)

Figure 7: Eigenvalues with bootstrap ranges, estimated subspace errors, and sufficient summary plots for SAVE (Algorithm 2) applied to B_{ind} from (59).

5. Summary and conclusion

We investigate sufficient dimension reduction as a tool for subspace-based dimension reduction in deterministic functions. We show that SDR may be used for ridge recovery by proving equivalence of the dimension reduction subspace

and the ridge subspace given $y = f(\mathbf{x})$. We define the central ridge subspace—the analog the central DRS—as the goal of the ridge recovery problem and conclude that SDR provides a theoretical framework for methods that search for the central ridge subspace.

We study the application of two SDR algorithms to the ridge recovery problem: sliced inverse regression and sliced average variance estimation. In regression, these methods use the inverse regression $\mathbf{x}|y$ to estimate subspaces relating to the central DRS. In ridge recovery, we reinterpret these methods as estimating integrals over the inverse image $f^{-1}(y)$, and we show that the SIR and SAVE subspaces are contained within the central ridge subspace. We also study the numerical convergence of the algorithms. The mean-squared error of the estimated SIR and SAVE eigenvalues converges at a rate inversely proportional to $N_{r_{\min}}$—the smallest number of samples in any slice. The distance between the estimated subspace and the true subspace decays like $N_{r_{\min}}^{-1}$ for both SIR and SAVE.

This paper's focus is ridge recovery via SDR; however, functions exhibiting ridge-like structure are present in a wide array of problems. SDR and inverse regression methods could provide be a valuable dimension reduction tool and provide insight into complex computer simulations. Understanding the role of SDR in ridge approximation requires thorough study of the eigenpairs of the SIR and SAVE matrices and their relationship to the structure of $f(\mathbf{x})$.

Appendix A. Proofs

We begin by introducing some notation and a short lemma which will appear in several proofs in this section. Define the multi-index set

$$\mathcal{K}^{p,q}(r) \;=\; \left\{ \mathbf{k} \in \mathbb{N}^{p+q} \;\middle|\; \begin{array}{c} k_1, \ldots k_p \in \{1, \ldots, N\} \\ \text{and } k_{p+1}, \ldots, k_{p+q} \in \{1, \ldots, N_r\} \end{array} \right\} \tag{A.1}$$

where $N_r \leq N$. Let $\mathbf{b} \in \mathbb{R}^{p+q}$ be a random vector and let $\left\{\mathbf{b}^k\right\}$, $k = 1, \ldots, N$, denote N independently drawn realizations of $\mathbf{b}$ where N is from (A.1). Define

the $(p+q)$-dimensional tensor $\boldsymbol{B}^{p,q}(r)$ whose elements are given by

$$\boldsymbol{B}^{p,q}_{\mathbf{k}}(r) \;=\; \mathbb{E}\left[b_1^{k_1}\dots b_p^{k_p} b_{p+1}^{k_{p+1}}\dots b_{p+q}^{k_{p+q}}\right] \tag{A.2}$$

for $\mathbf{k} \in \mathcal{K}^{p,q}(r)$. Note that the k_i's above do not indicate powers of b_j, but rather indices. That is, b_j^k denotes the jth entry of $\mathbf{b}^k$ which is the kth realization of the random vector $\mathbf{b}$. Thus, $\boldsymbol{B}^{p,q}(r)$ has dimensions

$$\boldsymbol{B}^{p,q}(r) \in \mathbb{R}^{\overbrace{N\times\cdots\times N}^{p}\times\overbrace{N_r\times\cdots\times N_r}^{q}}. \tag{A.3}$$

The following lemma plays an important role in upcoming proofs. This lemma is proved in section Appendix A.1 Let $\mathbf{b} \in \mathbb{R}^{p+q}$ be a random vector and let $\left\{\mathbf{b}^k\right\}$, $k = 1,\dots,N$, denote N independently drawn realizations of $\mathbf{b}$. Define $N_r \in \mathbb{N}$ such that $p+q \le N_r \le N$. Let $\mathcal{K}^{p,q}(r)$ and $\boldsymbol{B}^{p,q}(r)$ be defined according to (A.1) and (A.2), respectively. Then,

$$\sum_{\mathbf{k}\in\mathcal{K}^{p,q}(r)} \boldsymbol{B}^{p,q}_{\mathbf{k}}(r) \;=\; N^p N_r^q\,\mathbb{E}\left[b_1\right]\dots\mathbb{E}\left[b_p\right]\mathbb{E}\left[b_{p+1}\right]\dots\mathbb{E}\left[b_{p+q}\right] + \mathcal{O}(N^p N_r^{q-1} + N^{p-1}N_r^q). \tag{A.4}$$

Appendix A.1. Proof of Lemma Appendix A

Let $\mathbf{b} \in \mathbb{R}^{p+q}$ be a random vector and let $\left\{\mathbf{b}^k\right\}$, $k = 1,\dots,N$, denote N independently drawn realizations of $\mathbf{b}$. Define $N_r \in \mathbb{N}$ such that $p+q \le N_r \le N$. Let $\mathcal{K}^{p,q}(r)$ and $\boldsymbol{B}^{p,q}(r)$ be defined according to (A.1) and (A.2), respectively. Then,

$$\sum_{\mathbf{k}\in\mathcal{K}^{p,q}(r)} \boldsymbol{B}^{p,q}_{\mathbf{k}}(r) \;=\; N^p N_r^q\,\mathbb{E}\left[b_1\right]\dots\mathbb{E}\left[b_p\right]\mathbb{E}\left[b_{p+1}\right]\dots\mathbb{E}\left[b_{p+q}\right] + \mathcal{O}(N^p N_r^{q-1} + N^{p-1}N_r^q). \tag{A.4}$$

Proof. Define the following subsets of $\mathcal{K}^{p,q}(r)$:

$$\mathcal{K}_1^{p,q}(r) \;=\; \{\mathbf{k}\in\mathcal{K}^{p,q}(r)\,|k_i \neq k_j \text{ for } i,j = 1,\dots,p+q \text{ and } i\neq j\} \tag{A.5}$$

and $\mathcal{K}_2^{p,q}(r) = \mathcal{K}^{p,q}(r) \setminus \mathcal{K}_1^{p,q}(r)$. Asymptotically, as N and N_r tend towards infinity, the subset $\mathcal{K}_1^{p,q}(r)$ has

$$\left(\prod_{i=0}^{q-1}(N_r - i)\right)\left(\prod_{j=q}^{p+q-1}(N - j)\right) \;=\; N^p N_r^q + \mathcal{O}(N^p N_r^{q-1} + N^{p-1}N_r^q) \quad \text{(A.6)}$$

elements. Since all of the indices are unique for any $\mathbf{k} \in \mathcal{K}_1^{p,q}(r)$,

$$\boldsymbol{B}_{\mathbf{k}}^{p,q}(r) \;=\; \mathbb{E}\left[b_1\right] \ldots \mathbb{E}\left[b_p\right] \mathbb{E}\left[b_{p+1}\right] \ldots \mathbb{E}\left[b_{p+q}\right]. \quad \text{(A.7)}$$

Consider the number of elements in $\mathcal{K}_2^{p,q}(r)$. By construction, any $\mathbf{k} \in \mathcal{K}_2^{p,q}(r)$ must have at least two identical indices. As N and N_r tend towards infinity, the largest subset of elements in $\mathcal{K}_2^{p,q}(r)$ will be those with only two identical elements. There are three such cases:

Case 1:

Two of the k_i's which range from $1, \ldots, N$ are identical and all other k_i's are unique. There are

$$\left(\prod_{i=0}^{q-1}(N_r - i)\right)\left(\prod_{j=q}^{p+q-2}(N - j)\right) \;=\; \mathcal{O}(N^{p-1}N_r^q) \quad \text{(A.8)}$$

such elements.

Case 2:

Two of the k_i's which range from $1, \ldots, N_r$ are identical and all other k_i's are unique. There are

$$\left(\prod_{i=0}^{q-2}(N_r - i)\right)\left(\prod_{j=q-1}^{p+q-2}(N - j)\right) \;=\; \mathcal{O}(N^p N_r^{q-1}) \quad \text{(A.9)}$$

such elements.

Case 3:

One k_i which range from $1, \ldots, N$ and one k_i which ranges from $1, \ldots, N_r$ are identical and all other k_i's are unique. There are

$$\left(\prod_{i=0}^{q-1}(N_r - i)\right)\left(\prod_{j=q}^{p+q-2}(N - j)\right) \;=\; \mathcal{O}(N^{p-1}N_r^q) \quad \text{(A.10)}$$

such elements. Thus, there are $\mathcal{O}(N^p N_r^{q-1} + N^{p-1} N_r^q)$ elements in $\mathcal{K}_2^{p,q}(r)$.

Therefore,

$$\begin{aligned}
\sum_{\mathbf{k} \in \mathcal{K}^{p,q}(r)} \boldsymbol{B}_{\mathbf{k}}^{p,q}(r) &= \sum_{\mathbf{k} \in \mathcal{K}_1^{p,q}(r)} \boldsymbol{B}_{\mathbf{k}}^{p,q}(r) + \sum_{\mathbf{k} \in \mathcal{K}_2^{p,q}(r)} \boldsymbol{B}_{\mathbf{k}}^{p,q}(r) \\
&= \left(N^p N_r^q + \mathcal{O}(N^p N_r^{q-1} + N^{p-1} N_r^q)\right) \mathbb{E}\left[b_1\right] \ldots \mathbb{E}\left[b_p\right] \mathbb{E}\left[b_{p+1}\right] \ldots \mathbb{E}\left[b_{p+q}\right] \\
&\qquad + \mathcal{O}(N^p N_r^{q-1} + N^{p-1} N_r^q) \\
&= N^p N_r^q \, \mathbb{E}\left[b_1\right] \ldots \mathbb{E}\left[b_p\right] \mathbb{E}\left[b_{p+1}\right] \ldots \mathbb{E}\left[b_{p+q}\right] + \mathcal{O}(N^p N_r^{q-1} + N^{p-1} N_r^q)
\end{aligned} \tag{A.11}$$

□

Appendix A.2. Proof of Theorem 6

Theorem 6. *Let (Ω, Σ, P) be a probability triple with sample space Ω, σ-algebra Σ, and probability measure P. Suppose that $\mathbf{x} : \Omega \to \mathbb{R}^m$ and $y : \Omega \to \mathbb{R}$ are random variables related by a measurable function $f : \mathbb{R}^m \to \mathbb{R}$ so that $y = f(\mathbf{x})$. Let $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ be a constant matrix. Then $y \perp\!\!\!\perp \mathbf{x} | \boldsymbol{A}^T \mathbf{x}$ if and only if $y = g(\boldsymbol{A}^T \mathbf{x})$ where $g : \mathbb{R}^n \to \mathbb{R}$ is a measurable function.*

Proof. Assume that $y = f(\mathbf{x}) = g(\boldsymbol{A}^T \mathbf{x})$ for some $\boldsymbol{A}$ and g, then conditional independence follows since the value of $\boldsymbol{A}^T \mathbf{x}$ fully determines y.

Next, assume that $y \perp\!\!\!\perp \mathbf{x} | \boldsymbol{A}^T \mathbf{x}$, and note that $\boldsymbol{A}^T \mathbf{x}$ is a random vector from Ω to $\mathbb{R}^n$. The conditional independence $y \perp\!\!\!\perp \mathbf{x} | \boldsymbol{A}^T \mathbf{x}$ may be considered in terms of the σ-algebras generated by these random variables,

$$\sigma(y) \perp\!\!\!\perp \sigma(\mathbf{x}) | \sigma(\boldsymbol{A}^T \mathbf{x}) \tag{A.12}$$

where

$$\begin{aligned}
\sigma(y) &= \left\{ y^{-1}(B) \,:\, B \in \mathcal{B}(\mathbb{R}) \right\}, \\
\sigma(\mathbf{x}) &= \left\{ \mathbf{x}^{-1}(B) \,:\, B \in \mathcal{B}(\mathbb{R}^m) \right\}, \\
\sigma(\boldsymbol{A}^T \mathbf{x}) &= \left\{ (\boldsymbol{A}^T \mathbf{x})^{-1}(B) \,:\, B \in \mathcal{B}(\mathbb{R}^n) \right\},
\end{aligned} \tag{A.13}$$

and $\mathcal{B}(\cdot)$ is the Borel σ-algebra over the indicated domain. Equation (A.12) is equivalent to

$$\mathbb{E}\left[z \,\middle|\, \sigma(\boldsymbol{A}^T \mathbf{x}) \right] = \mathbb{E}\left[z \,\middle|\, \sigma\left(\sigma(\boldsymbol{A}^T \mathbf{x}) \cup \sigma(\mathbf{x}) \right) \right] \tag{A.14}$$

where $z : \Omega \to \mathbb{R}^+$ is any $\sigma(y)$-measurable random variable [34, Proposition 13, Chapter 3].

The function $y = f(\mathbf{x})$ may be partitioned into its positive and negative components such that $y = y^+ - y^-$ where

$$\begin{aligned} y^+ &= f^+(\mathbf{x}) = \begin{cases} f(\mathbf{x}) & \text{if } f(\mathbf{x}) > 0 \\ 0 & \text{otherwise} \end{cases}, \\ y^- &= f^-(\mathbf{x}) = \begin{cases} -f(\mathbf{x}) & \text{if } f(\mathbf{x}) < 0 \\ 0 & \text{otherwise} \end{cases}. \end{aligned} \tag{A.15}$$

Since y is measurable, $y^+ : \Omega \to \mathbb{R}^+$ is $\sigma(y)$-measurable [35]. Applying (A.14) to y^+,

$$\mathbb{E}\left[y^+ \,\middle|\, \sigma(\boldsymbol{A}^T\mathbf{x})\right] = \mathbb{E}\left[y^+ \,\middle|\, \sigma\left(\sigma(\boldsymbol{A}^T\mathbf{x}) \cup \sigma(\mathbf{x})\right)\right]. \tag{A.16}$$

Consider $\sigma(\boldsymbol{A}^T\mathbf{x})$. We can view $\boldsymbol{A}^T\mathbf{x}$ as a function $h : \mathbb{R}^m \to \mathbb{R}^n$ which maps $\mathbf{x}$ to $\boldsymbol{A}^T\mathbf{x}$. Since h is a linear function, it is $\mathcal{B}(\mathbb{R}^n)$-measurable so that $h^{-1}(B) \in \mathcal{B}(\mathbb{R}^m)$ for any $B \in \mathcal{B}(\mathbb{R}^n)$. This implies that $\sigma(h(\mathbf{x})) = \sigma(\boldsymbol{A}^T\mathbf{x}) \subseteq \sigma(\mathbf{x})$. Therefore, $\sigma(\boldsymbol{A}^T\mathbf{x}) \cup \sigma(\mathbf{x}) = \sigma(\mathbf{x})$, and

$$\mathbb{E}\left[y^+ \,\middle|\, \sigma(\boldsymbol{A}^T\mathbf{x})\right] = \mathbb{E}\left[y^+ \,|\, \sigma(\mathbf{x})\right]. \tag{A.17}$$

The Doob-Dynkin Lemma [34, Proposition 3, Chapter 1] states that, since the deterministic relationship f^+ exists between y^+ and $\mathbf{x}$, y^+ is $\sigma(\mathbf{x})$-measurable. Therefore,

$$y^+ = f^+(\mathbf{x}) = \mathbb{E}\left[y^+ \,|\sigma(\mathbf{x})\right] = \mathbb{E}\left[y^+ \,\middle|\, \sigma(\boldsymbol{A}^T\mathbf{x})\right] \tag{A.18}$$

by (A.17). The above argument also applies for $y^- : \Omega \to \mathbb{R}^+$ so that

$$y^- = f^-(\mathbf{x}) = \mathbb{E}\left[y^- \,|\, \sigma(\mathbf{x})\right] = \mathbb{E}\left[y^- \,\middle|\, \sigma(\boldsymbol{A}^T\mathbf{x})\right]. \tag{A.19}$$

Thus,

$$\begin{aligned} y &= y^+ - y^- \\ &= \mathbb{E}\left[y^+ \,\middle|\, \sigma(\boldsymbol{A}^T\mathbf{x})\right] - \mathbb{E}\left[y^- \,\middle|\, \sigma(\boldsymbol{A}^T\mathbf{x})\right] \\ &= \mathbb{E}\left[y^+ - y^- \,\middle|\, \sigma(\boldsymbol{A}^T\mathbf{x})\right] \\ &= \mathbb{E}\left[y \,\middle|\, \sigma(\boldsymbol{A}^T\mathbf{x})\right]. \end{aligned} \tag{A.20}$$

Consider $\mathbb{E}\left[y|\sigma(\boldsymbol{A}^T\mathbf{x})\right]$ which is the restriction of $f(\mathbf{x})$ to $\sigma(\boldsymbol{A}^T\mathbf{x})$. Thus, y is $\sigma(\boldsymbol{A}^T\mathbf{x})$-measurable, and by the Doob-Dynkin Lemma [34, Proposition 3, Chapter 1], it is a function of $\boldsymbol{A}^T\mathbf{x}$. That is,

$$\mathbb{E}\left[y|\sigma(\boldsymbol{A}^T\mathbf{x})\right] \;=\; g(\boldsymbol{A}^T\mathbf{x}) \tag{A.21}$$

where $g : \mathbb{R}^n \to \mathbb{R}$ is a deterministic measurable function. Combining (A.20) and (A.21),

$$y \;=\; g(\boldsymbol{A}^T\mathbf{x}). \tag{A.22}$$

□

Appendix A.3. Proof of Lemma 3

Let $y = f(\mathbf{x})$ with $\mathbf{x} \in \mathbb{R}^m$ and $y \in \mathbb{R}$. Suppose that $f(\mathbf{x})$ may be written $y = f(\mathbf{x}) = g_1(\boldsymbol{A}_1^T\mathbf{x})$ for some $\boldsymbol{A}_1 \in \mathbb{R}^{m\times n_1}$ and $g_1 : \mathbb{R}^{n_1} \to \mathbb{R}$. If $\boldsymbol{A}_2 \in \mathbb{R}^{m\times n_2}$ is such that $\text{colspan}(\boldsymbol{A}_1) = \text{colspan}(\boldsymbol{A}_2)$, then there exists $g_2 : \mathbb{R}^{n_2} \to \mathbb{R}$ such that

$$y \;=\; f(\mathbf{x}) \;=\; g_2(\boldsymbol{A}_2^T\mathbf{x}). \tag{26}$$

Proof. Since $\text{colspan}(\boldsymbol{A}_1) = \text{colspan}(\boldsymbol{A}_2)$, each column of $\boldsymbol{A}_1$ can be expressed as a linear combination of columns of $\boldsymbol{A}_2$,

$$\boldsymbol{A}_1 \;=\; \boldsymbol{A}_2\boldsymbol{B}. \tag{A.23}$$

Then,

$$y \;=\; f(\mathbf{x}) \;=\; g_1(\boldsymbol{A}_1^T\mathbf{x}) \;=\; g_1(\boldsymbol{B}^T\boldsymbol{A}_2^T\mathbf{x}) \;=\; g_2(\boldsymbol{A}_2^T\mathbf{x}) \tag{A.24}$$

where $g_2 : \mathbb{R}^{n_2} \to \mathbb{R}$. □

Appendix A.4. Proof of Lemma 3

Assume that the support of the input density $\rho(\mathbf{x})$ is convex. Let $y = f(\mathbf{x})$, and assume f is differentiable for all $\mathbf{x}$ in the support of $\rho(\mathbf{x})$. Suppose that $\mathcal{S}_1$ and $\mathcal{S}_2$ are two ridge subspaces for $f(\mathbf{x})$. Then, $\mathcal{S}_1 \cap \mathcal{S}_2$ is also a ridge subspace.

Proof. Without loss of generality, let $\boldsymbol{A}_1 \in \mathbb{R}^{m \times n_1}$ and $\boldsymbol{A}_2 \in \mathbb{R}^{m \times n_2}$ contain a linearly independent basis for $\mathcal{S}_1$ and $\mathcal{S}_2$, respectively, such that

$$\boldsymbol{A}_1 \;=\; \begin{bmatrix} \boldsymbol{A} & \bar{\boldsymbol{A}}_1 \end{bmatrix}, \qquad \boldsymbol{A}_2 \;=\; \begin{bmatrix} \boldsymbol{A} & \bar{\boldsymbol{A}}_2 \end{bmatrix} \tag{A.25}$$

where $\operatorname{colspan}(\bar{\boldsymbol{A}}_1) = \mathcal{S}_1 \backslash (\mathcal{S}_2 \cap \mathcal{S}_1)$, $\operatorname{colspan}(\bar{\boldsymbol{A}}_2) = \mathcal{S}_2 \backslash (\mathcal{S}_1 \cap \mathcal{S}_2)$, and $\operatorname{colspan}(\boldsymbol{A}) = \mathcal{S}_1 \cap \mathcal{S}_2$. Define the following variables:

$$\mathbf{u}_1 \;=\; \boldsymbol{A}_1^T \mathbf{x}, \qquad \mathbf{u}_2 \;=\; \boldsymbol{A}_2^T \mathbf{x}, \qquad \mathbf{u} \;=\; \boldsymbol{A}^T \mathbf{x}, \qquad \bar{\mathbf{u}}_1 \;=\; \bar{\boldsymbol{A}}_1^T \mathbf{x}, \qquad \bar{\mathbf{u}}_2 \;=\; \bar{\boldsymbol{A}}_2^T \mathbf{x}. \tag{A.26}$$

Let $g_1 : \mathbb{R}^{n_1} \to \mathbb{R}$ and $g_2 : \mathbb{R}^{n_2} \to \mathbb{R}$ denote the ridge profiles associated with $\mathcal{S}_1$ and $\mathcal{S}_2$, respectively. Then,

$$y \;=\; f(\mathbf{x}) \;=\; g_1(\boldsymbol{A}_1^T \mathbf{x}) \;=\; g_1(\mathbf{u}_1) \;=\; g_2(\boldsymbol{A}_2^T \mathbf{x}) \;=\; g_2(\mathbf{u}_2). \tag{A.27}$$

Substituting (A.25) into (A.27),

$$\begin{aligned} y \;&=\; f(\mathbf{x}) \\ &=\; g_1\left(\begin{bmatrix} \boldsymbol{A}^T \\ \bar{\boldsymbol{A}}_1^T \end{bmatrix} \mathbf{x}\right) \;=\; g_1\left(\begin{bmatrix} \mathbf{u} \\ \bar{\mathbf{u}}_1 \end{bmatrix}\right) \\ &=\; g_2\left(\begin{bmatrix} \boldsymbol{A}^T \\ \bar{\boldsymbol{A}}_2^T \end{bmatrix} \mathbf{x}\right) \;=\; g_2\left(\begin{bmatrix} \mathbf{u} \\ \bar{\mathbf{u}}_2 \end{bmatrix}\right). \end{aligned} \tag{A.28}$$

Taking the gradient of (A.28) with respect to $\mathbf{x}$,

$$\begin{aligned} \nabla_{\mathbf{x}} y \;&=\; \nabla_{\mathbf{x}} f(\mathbf{x}) \\ &=\; \nabla_{\mathbf{x}} g_1\left(\begin{bmatrix} \boldsymbol{A}^T \\ \bar{\boldsymbol{A}}_1^T \end{bmatrix} \mathbf{x}\right) \;=\; \boldsymbol{A} \nabla_{\mathbf{u}} g_1\left(\begin{bmatrix} \mathbf{u} \\ \bar{\mathbf{u}}_1 \end{bmatrix}\right) + \bar{\boldsymbol{A}}_1 \nabla_{\bar{\mathbf{u}}_1} g_1\left(\begin{bmatrix} \mathbf{u} \\ \bar{\mathbf{u}}_1 \end{bmatrix}\right) \\ &=\; \nabla_{\mathbf{x}} g_2\left(\begin{bmatrix} \boldsymbol{A}^T \\ \bar{\boldsymbol{A}}_2^T \end{bmatrix} \mathbf{x}\right) \;=\; \boldsymbol{A} \nabla_{\mathbf{u}} g_2\left(\begin{bmatrix} \mathbf{u} \\ \bar{\mathbf{u}}_2 \end{bmatrix}\right) + \bar{\boldsymbol{A}}_2 \nabla_{\bar{\mathbf{u}}_2} g_2\left(\begin{bmatrix} \mathbf{u} \\ \bar{\mathbf{u}}_2 \end{bmatrix}\right). \end{aligned} \tag{A.29}$$

Then,

$$\begin{aligned}
\boldsymbol{A}\nabla_{\mathbf{u}} g_1\left(\begin{bmatrix}\mathbf{u}\\ \bar{\mathbf{u}}_1\end{bmatrix}\right) + \bar{\boldsymbol{A}}_1\nabla_{\bar{\mathbf{u}}_1} g_1\left(\begin{bmatrix}\mathbf{u}\\ \bar{\mathbf{u}}_1\end{bmatrix}\right) &= \boldsymbol{A}\nabla_{\mathbf{u}} g_2\left(\begin{bmatrix}\mathbf{u}\\ \bar{\mathbf{u}}_2\end{bmatrix}\right) + \bar{\boldsymbol{A}}_2\nabla_{\bar{\mathbf{u}}_2} g_2\left(\begin{bmatrix}\mathbf{u}\\ \bar{\mathbf{u}}_2\end{bmatrix}\right)\\
\boldsymbol{A}\nabla_{\mathbf{u}}\left(g_1\left(\begin{bmatrix}\mathbf{u}\\ \bar{\mathbf{u}}_1\end{bmatrix}\right) - g_2\left(\begin{bmatrix}\mathbf{u}\\ \bar{\mathbf{u}}_2\end{bmatrix}\right)\right) + \bar{\boldsymbol{A}}_1\nabla_{\bar{\mathbf{u}}_1} g_1\left(\begin{bmatrix}\mathbf{u}\\ \bar{\mathbf{u}}_1\end{bmatrix}\right) &= \bar{\boldsymbol{A}}_2\nabla_{\bar{\mathbf{u}}_2} g_2\left(\begin{bmatrix}\mathbf{u}\\ \bar{\mathbf{u}}_2\end{bmatrix}\right)\\
\bar{\boldsymbol{A}}_1\nabla_{\bar{\mathbf{u}}_1} g_1\left(\begin{bmatrix}\mathbf{u}\\ \bar{\mathbf{u}}_1\end{bmatrix}\right) &= \bar{\boldsymbol{A}}_2\nabla_{\bar{\mathbf{u}}_2} g_2\left(\begin{bmatrix}\mathbf{u}\\ \bar{\mathbf{u}}_2\end{bmatrix}\right).
\end{aligned} \tag{A.30}$$

The left-hand side of (A.30) is an element of $\text{colspan}(\bar{\boldsymbol{A}}_1)$ and the right-hand side is an element of $\text{colspan}(\bar{\boldsymbol{A}}_2)$. However, by construction $\text{colspan}(\bar{\boldsymbol{A}}_1) \cap \text{colspan}(\bar{\boldsymbol{A}}_2) = \{\mathbf{0}\}$. Hence,

$$\bar{\boldsymbol{A}}_1\nabla_{\bar{\mathbf{u}}_1} g_1\left(\begin{bmatrix}\mathbf{u}\\ \bar{\mathbf{u}}_1\end{bmatrix}\right) = \bar{\boldsymbol{A}}_2\nabla_{\bar{\mathbf{u}}_2} g_2\left(\begin{bmatrix}\mathbf{u}\\ \bar{\mathbf{u}}_2\end{bmatrix}\right) = \mathbf{0}. \tag{A.31}$$

Thus,

$$\nabla_{\bar{\mathbf{u}}_1} g_1\left(\begin{bmatrix}\mathbf{u}\\ \bar{\mathbf{u}}_1\end{bmatrix}\right) = \nabla_{\bar{\mathbf{u}}_2} g_2\left(\begin{bmatrix}\mathbf{u}\\ \bar{\mathbf{u}}_2\end{bmatrix}\right) = \mathbf{0} \tag{A.32}$$

which implies that g_1 and g_2 are constant along $\bar{\mathbf{u}}_1$ and $\bar{\mathbf{u}}_2$, respectively. There exists a ridge profile $g : \mathbb{R}^n \to \mathbb{R}$ such that

$$y = f(\mathbf{x}) = g(\boldsymbol{A}^T\mathbf{x}) = g(\mathbf{u}). \tag{A.33}$$

Then, $\mathcal{S} = \mathcal{S}_1 \cap \mathcal{S}_2 = \text{colspan}(\boldsymbol{A})$ defines a ridge subspace for $y = f(\mathbf{x})$. □

Appendix A.5. Proof of Property 3

Property 3. *Let* $\mathbf{z} = \boldsymbol{B}\mathbf{x} + \mathbf{b}$ *be an affine transformation of inputs* $\mathbf{x}$*, where* $\boldsymbol{B} \in \mathbb{R}^{m\times m}$ *is full rank and* $\mathbf{b} \in \mathbb{R}^m$*. Given* $y = f(\mathbf{x})$*, define*

$$\tilde{f}(\mathbf{z}) = f\left(\boldsymbol{B}^{-1}(\mathbf{z} - \mathbf{b})\right). \tag{27}$$

Let $\mathcal{S}_{f(\mathbf{x})}$ *denote the central ridge subspace for* $f(\mathbf{x})$ *with basis matrix* $\boldsymbol{A}$*. Then,* $\text{colspan}(\boldsymbol{B}^{-T}\boldsymbol{A})$ *is the central ridge subspace for* $\tilde{f}(\mathbf{z})$*.*

Proof. Let g denote the ridge profile of $f(\mathbf{x})$. Then,

$$\begin{aligned} y \;=\; f(\mathbf{x}) \;&=\; g\left(\boldsymbol{A}^T\mathbf{x}\right) \\ &=\; g\left(\boldsymbol{A}^T\left(\boldsymbol{B}^{-1}(\mathbf{z}-\mathbf{b})\right)\right) \\ &=\; g\left(\boldsymbol{A}^T\boldsymbol{B}^{-1}\mathbf{z}-\boldsymbol{A}^T\boldsymbol{B}^{-1}\mathbf{b}\right) \\ &=\; \tilde{g}\left(\boldsymbol{A}^T\boldsymbol{B}^{-1}\mathbf{z}\right) \\ &=\; \tilde{g}\left(\left(\boldsymbol{B}^{-T}\boldsymbol{A}\right)^T\mathbf{z}\right). \end{aligned} \tag{A.34}$$

Therefore, the matrix $\boldsymbol{B}^{-T}\boldsymbol{A} \in \mathbb{R}^{m\times n}$ contains a basis for the central ridge subspace of $\tilde{f}(\mathbf{z})$. □

Appendix A.6. Proof of Property 4

Property 4. *Let $\mathcal{S}_{f(\mathbf{x})}$ be the central ridge subspace for $y = f(\mathbf{x})$ and $h(y)$ be any function of the output. Then,*

$$\mathcal{S}_{h(f(\mathbf{x}))} \subseteq \mathcal{S}_{f(\mathbf{x})}. \tag{29}$$

Proof. Consider $h(y)$,

$$h(y) = h(f(\mathbf{x})) = h(g(\boldsymbol{A}^T\mathbf{x})) = \tilde{g}(\boldsymbol{A}^T\mathbf{x}). \tag{A.35}$$

Thus, $h(y)$ is a ridge function with $\mathcal{S}_{h(f(\mathbf{x}))} \subseteq \mathcal{S}_{f(\mathbf{x})}$.

To show a strict subset, consider $h(y) = k$ where k is a constant. From section 3, $\mathcal{S}_{h(f(\mathbf{x}))} = \{\mathbf{0}\}$. Choosing nonconstant $f(\mathbf{x})$, we get $\mathcal{S}_{h(f(\mathbf{x}))} = \{\mathbf{0}\}$ to be a strict subspace of $\mathcal{S}_{f(\mathbf{x})}$. □

Appendix A.7. Proof of Lemma 3.1

A vector $\mathbf{w} \in \mathbb{R}^m$ is in the null space of $\boldsymbol{C}_{\text{SIR}}$ if and only if $\boldsymbol{\mu}(y)^T\mathbf{w} = 0$ for all y in the range of f.

Proof. For any $\mathbf{w} \in \mathbb{R}^m$,

$$\mathbf{w}^T\boldsymbol{C}_{\text{SIR}}\mathbf{w} \;=\; \mathbf{w}^T\left(\int \boldsymbol{\mu}(y)\,\boldsymbol{\mu}(y)^T\,d\gamma(y)\right)\mathbf{w} \;=\; \int \left[\boldsymbol{\mu}(y)^T\mathbf{w}\right]^2\,d\gamma(y). \tag{A.36}$$

Assume $\mathbf{w} \in \text{null}(\boldsymbol{C}_{\text{SIR}})$. Then, $\mathbf{w}^T \boldsymbol{C}_{\text{SIR}} \mathbf{w} = 0$ which implies

$$\int \left[\boldsymbol{\mu}(y)^T \mathbf{w}\right]^2 \, d\gamma(y) = 0. \tag{A.37}$$

Since $\left[\boldsymbol{\mu}(y)^T \mathbf{w}\right]^2 \geq 0$ and γ is a measure on the range of y, we have $\boldsymbol{\mu}(y)^T \mathbf{w} = 0$.

Next, assume $\boldsymbol{\mu}(y)^T \mathbf{w} = 0$ for all y in the range of f. Then,

$$\int \left[\boldsymbol{\mu}(y)^T \mathbf{w}\right]^2 \, d\gamma(y) = 0. \tag{A.38}$$

This implies $\mathbf{w}^T \boldsymbol{C}_{\text{SIR}} \mathbf{w} = 0$, or equivalently $\mathbf{w} \in \text{null}(\boldsymbol{C}_{\text{SIR}})$. □

Appendix A.8. Proof of Theorem 7

Theorem 7. *Given an elliptically symmetric input probability density function $\rho : \mathbb{R}^m \to \mathbb{R}^+$, assume $f : \mathbb{R}^m \to \mathbb{R}$ admits a central ridge subspace $\mathcal{S}_{f(\mathbf{x})}$. Then,* $\text{colspan}(\boldsymbol{C}_{\text{SIR}}) \subseteq \mathcal{S}_{f(\mathbf{x})}$.

Proof. Consider the orthogonal $m \times m$ matrix $\begin{bmatrix} \boldsymbol{A} & \boldsymbol{B} \end{bmatrix}$ where $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ contains a basis for $\mathcal{S}_{f(\mathbf{x})}$ and $\boldsymbol{B} \in \mathbb{R}^{m \times (m-n)}$ is a basis for the orthogonal complement of $\mathcal{S}_{f(\mathbf{x})}$. Define the following decomposition of $\mathbf{x}$:

$$\mathbf{x} \;=\; \boldsymbol{A}\mathbf{u} + \boldsymbol{B}\mathbf{v} \tag{A.39}$$

where $\mathbf{u} = \boldsymbol{A}^T \mathbf{x} \in \mathbb{R}^n$ and $\mathbf{v} = \boldsymbol{B}^T \mathbf{x} \in \mathbb{R}^{m-n}$.

By (33),

$$\boldsymbol{\mu}(y) \;=\; \int \left[\int \mathbf{x} \, d\mu_{\mathbf{x}|\mathbf{u}}(\mathbf{x}) \right] d\mu_{\mathbf{u}|y}(\mathbf{u}). \tag{A.40}$$

Consider the inner integral of (A.40). The conditional measure $\mu_{\mathbf{x}|\mathbf{u}}$ is over the set $\mathcal{X}(\mathbf{u}) = \left\{ \mathbf{x} \in \mathbb{R}^m \,|\, \mathbf{u} = \boldsymbol{A}^T \mathbf{x} \right\}$. Applying the decomposition of $\mathbf{x}$ from (A.39) to $\mathcal{X}(\mathbf{u})$ for a fixed $\mathbf{u}^*$ yields

$$\begin{aligned} \mathcal{X}(\mathbf{u}^*) \;&=\; \left\{ \mathbf{x} \in \mathbb{R}^m \,|\, \mathbf{u}^* = \boldsymbol{A}^T \mathbf{x} \right\} \\ &=\; \left\{ \mathbf{x} \in \mathbb{R}^m \,|\, \mathbf{x} = \boldsymbol{A}\mathbf{u}^* + \boldsymbol{B}\mathbf{v},\, \mathbf{v} \in \mathbb{R}^{m-n} \right\}. \end{aligned} \tag{A.41}$$

Let $\nu_{\mathbf{v}|\mathbf{u}}$ denote the conditional measure on $\mathbf{v} \in \mathbb{R}^{m-n}$ induced by $\mathbf{v} = \boldsymbol{B}^T \mathbf{x}$ and $\rho(\mathbf{x})$ over the set $\mathcal{X}(\mathbf{u})$. Since $\mathbf{x}$ is standardized with an elliptically symmetric density and $\begin{bmatrix} \boldsymbol{A} & \boldsymbol{B} \end{bmatrix}$ is an orthogonal matrix, $\nu_{\mathbf{v}|\mathbf{u}}$ is standardized and elliptically symmetric [36, Chapter 6].

Using this measure, the inner integral of (A.40) may be written

$$\begin{aligned}\int \mathbf{x}\, d\mu_{\mathbf{x}|\mathbf{u}}(\mathbf{x}) &= \int (\boldsymbol{A}\mathbf{u} + \boldsymbol{B}\mathbf{v})\; d\nu_{\mathbf{v}|\mathbf{u}}(\mathbf{v}) \\ &= \boldsymbol{A}\mathbf{u} \int d\nu_{\mathbf{v}|\mathbf{u}}(\mathbf{v}) + \boldsymbol{B} \int \mathbf{v}\, d\nu_{\mathbf{v}|\mathbf{u}}(\mathbf{v}) \\ &= \boldsymbol{A}\mathbf{u}\end{aligned} \tag{A.42}$$

Plugging the above result into (A.40) yields

$$\begin{aligned}\boldsymbol{\mu}(y) &= \int \left[\int \mathbf{x}\, d\mu_{\mathbf{x}|\mathbf{u}}(\mathbf{x}) \right] d\mu_{\mathbf{u}|y}(\mathbf{u}) \\ &= \boldsymbol{A} \int \mathbf{u}\, d\mu_{\mathbf{u}|y}(\mathbf{u}).\end{aligned} \tag{A.43}$$

Let $\mathbf{w} \in \text{null}(\boldsymbol{A}^T)$. Then,

$$\begin{aligned}\boldsymbol{\mu}(y)^T \mathbf{w} &= \left(\boldsymbol{A} \int \mathbf{u}\, d\mu_{\mathbf{u}|y}(\mathbf{u}) \right)^T \mathbf{w} \\ &= \left(\int \mathbf{u}^T\, d\mu_{\mathbf{u}|y}(\mathbf{u}) \right) (\boldsymbol{A}^T \mathbf{w}) \\ &= 0.\end{aligned} \tag{A.44}$$

By Lemma 3.1, $\mathbf{w} \in \text{null}(\boldsymbol{C}_{\text{SIR}})$. Therefore, $\text{null}(\boldsymbol{A}^T) \subseteq \text{null}(\boldsymbol{C}_{\text{SIR}})$. Since $\mathcal{S}_{f(\mathbf{x})} = \text{colspan}(\boldsymbol{A})$ is the orthogonal complement of $\text{null}(\boldsymbol{A}^T)$,

$$\text{colspan}(\boldsymbol{C}_{\text{SIR}}) \subseteq \mathcal{S}_{f(\mathbf{x})} \tag{A.45}$$

as required. □

Appendix A.9. Proof of Theorem 8

Theorem 8. *Assume that Algorithm 1 has been applied to the data set* $\{[\, y_i\, ,\, \mathbf{x}_i^T\,]\}$*, with* $i = 1, \ldots, N$*, where the* $\mathbf{x}_i$ *are drawn independently according to* $\rho(\mathbf{x})$ *and* $y_i = f(\mathbf{x}_i)$ *are point evaluations of* $f(\mathbf{x})$*. Then, for* $k = 1, \ldots, m$*,*

$$\mathbb{E}\left[\left(\lambda_k(\boldsymbol{D}_{\text{SIR}}) - \lambda_k(\hat{\boldsymbol{D}}_{\text{SIR}}) \right)^2 \right] \;=\; \mathcal{O}(N_{r_{\min}}^{-1}) \tag{40}$$

where $\lambda_k(\cdot)$ *denotes the* k*th eigenvalue of the given matrix.*

Proof. Recall the SIR matrix from (35),

$$\boldsymbol{D}_{\text{SIR}} \;=\; \sum_{r=1}^{R} \omega(r) \left[\boldsymbol{\mu}(r)\boldsymbol{\mu}(r)^T\right] \tag{A.46}$$

where $\omega(r)$ is the probability mass function from (36) and $\boldsymbol{\mu}(r)$ is the average of the inverse image $f^{-1}(h^{-1}(r))$ from (37). Algorithm 1 approximates $\boldsymbol{D}_{\text{SIR}}$ using

$$\hat{\boldsymbol{D}}_{\text{SIR}} \;=\; \sum_{r=1}^{R} \hat{\omega}(r) \left[\hat{\boldsymbol{\mu}}(r)\hat{\boldsymbol{\mu}}(r)^T\right] \tag{A.47}$$

where $\hat{\omega}(r)$ and $\hat{\boldsymbol{\mu}}(r)$ are sample estimates of $\omega(r)$ and $\boldsymbol{\mu}(r)$, respectively,

$$\hat{\omega}(r) \;=\; \frac{N_r}{N} \qquad \text{and} \qquad \hat{\boldsymbol{\mu}}(r) \;=\; \frac{1}{N_r} \sum_{i \in \mathcal{I}_r} \mathbf{x}_i \tag{A.48}$$

where $\mathcal{I}_r$ is the set of indices for which $y_i \in J_r$ and N_r is the cardinality of $\mathcal{I}_r$. We may rewrite $\hat{\omega}(r)$ as

$$\hat{\omega}(r) \;=\; \frac{1}{N} \sum_{i=1}^{N} \chi\left(y_i \in J_r\right) \tag{A.49}$$

where $\chi\left(y_i \in J_r\right)$ is an indicator function which is 1 when $y_i \in J_r$ and 0 otherwise.

We want to examine the mean squared error of the elements of $\hat{\boldsymbol{D}}_{\text{SIR}}$. First, consider a single term in the summation from (A.47). Fix r and let

$$\omega \;=\; \omega(r), \quad \hat{\omega} \;=\; \hat{\omega}(r), \quad \chi_i \;=\; \chi\left(y_i \in J_r\right), \quad \boldsymbol{\mu} \;=\; \boldsymbol{\mu}(r), \quad \text{and} \quad \hat{\boldsymbol{\mu}} \;=\; \hat{\boldsymbol{\mu}}(r). \tag{A.50}$$

Assume without loss of generality that $\mathcal{I}_r = \{1, \ldots, N_r\}$. This simplifies notation as we compute the mean squared error.

To compute the mean squared error, we focus on the computation of a single element in $\hat{\boldsymbol{D}}_{\text{SIR}}$. To do this, we move the sample index to the superscript and let the subscript denote the vector or matrix element. That is, we let $\hat{\mathbf{x}}_i^k$ denote the ith element of the vector $\hat{\mathbf{x}}^k$ which is the kth realization of the random

vector $\mathbf{x}$. Thus, for $1 \leq i, j \leq m$,

$$\begin{aligned}\mathbb{E}\left[\hat{\omega}\hat{\boldsymbol{\mu}}_i\hat{\boldsymbol{\mu}}_j\right] &= \mathbb{E}\left[\left(\frac{1}{N}\sum_{k_1=1}^{N}\chi^{k_1}\right)\left(\frac{1}{N_r}\sum_{k_2=1}^{N_r}\mathbf{x}_i^{k_2}\right)\left(\frac{1}{N_r}\sum_{k_3=1}^{N_r}\mathbf{x}_j^{k_3}\right)\right] \\ &= \frac{1}{N\,N_r^2}\sum_{k_1=1}^{N}\sum_{k_2=1}^{N_r}\sum_{k_3=1}^{N_r}\mathbb{E}\left[\chi^{k_1}\,\mathbf{x}_i^{k_2}\,\mathbf{x}_j^{k_3}\right].\end{aligned} \tag{A.51}$$

Equations (A.1) and (A.2) allow us to rewrite (A.51) as a summation over a 3-dimensional tensor. Since we assume a fixed r, we drop the argument from the notation of these equations. Thus,

$$\mathbb{E}\left[\hat{\omega}\hat{\boldsymbol{\mu}}_i\hat{\boldsymbol{\mu}}_j\right] = \frac{1}{N\,N_r^2}\sum_{\mathbf{k}\in\mathcal{K}^{1,2}}\boldsymbol{B}_{\mathbf{k}}^{1,2}. \tag{A.52}$$

From Lemma Appendix A,

$$\begin{aligned}\mathbb{E}\left[\hat{\omega}\hat{\boldsymbol{\mu}}_i\hat{\boldsymbol{\mu}}_j\right] &= \frac{1}{N\,N_r^2}\sum_{\mathbf{k}\in\mathcal{K}^{1,2}}\boldsymbol{B}_{\mathbf{k}}^{1,2} \\ &= \frac{1}{N\,N_r^2}\left[N\,N_r^2\,\mathbb{E}\left[\chi\right]\mathbb{E}\left[\mathbf{x}_i\right]\mathbb{E}\left[\mathbf{x}_j\right] + \mathcal{O}(N\,N_r + N_r^2)\right] \\ &= \mathbb{E}\left[\chi\right]\mathbb{E}\left[\mathbf{x}_i\right]\mathbb{E}\left[\mathbf{x}_j\right] + \mathcal{O}(N_r^{-1}) \\ &= \omega\boldsymbol{\mu}_i\boldsymbol{\mu}_j + \mathcal{O}(N_r^{-1}).\end{aligned} \tag{A.53}$$

Next, we compute the variance using the property

$$\mathrm{Var}\left[\hat{\omega}\hat{\boldsymbol{\mu}}_i\hat{\boldsymbol{\mu}}_j\right] = \mathbb{E}\left[\left(\hat{\omega}\hat{\boldsymbol{\mu}}_i\hat{\boldsymbol{\mu}}_j\right)^2\right] - \mathbb{E}\left[\hat{\omega}\hat{\boldsymbol{\mu}}_i\hat{\boldsymbol{\mu}}_j\right]^2. \tag{A.54}$$

To find $\mathbb{E}\left[\left(\hat{\omega}\hat{\boldsymbol{\mu}}_i\hat{\boldsymbol{\mu}}_j\right)^2\right]$,

$$\begin{aligned}\mathbb{E}\left[\left(\hat{\omega}\hat{\boldsymbol{\mu}}_i\hat{\boldsymbol{\mu}}_j\right)^2\right] &= \mathbb{E}\left[\left(\frac{1}{N}\sum_{k_1=1}^{N}\chi^{k_1}\right)^2\left(\frac{1}{N_r}\sum_{k_2=1}^{N_r}\mathbf{x}_i^{k_2}\right)^2\left(\frac{1}{N_r}\sum_{k_3=1}^{N_r}\mathbf{x}_j^{k_3}\right)^2\right] \\ &= \frac{1}{N^2N_r^4}\sum_{k_1=1}^{N}\sum_{k_2=1}^{N}\sum_{k_3=1}^{N_r}\sum_{k_4=1}^{N_r}\sum_{k_5=1}^{N_r}\sum_{k_6=1}^{N_r}\mathbb{E}\left[\chi^{k_1}\chi^{k_2}\mathbf{x}_i^{k_3}\mathbf{x}_i^{k_4}\mathbf{x}_j^{k_5}\mathbf{x}_j^{k_6}\right].\end{aligned} \tag{A.55}$$

Again, we can express this summation using (A.1) and (A.2) as

$$\mathbb{E}\left[\left(\hat{\omega}\hat{\boldsymbol{\mu}}_i\hat{\boldsymbol{\mu}}_j\right)^2\right] = \frac{1}{N^2N_r^4}\sum_{\mathbf{k}\in\mathcal{K}^{2,4}}\boldsymbol{B}_{\mathbf{k}}^{2,4}. \tag{A.56}$$

By Lemma Appendix A,

$$\begin{aligned}
\mathbb{E}\left[(\hat{\omega}\hat{\boldsymbol{\mu}}_i\hat{\boldsymbol{\mu}}_j)^2\right] &= \frac{1}{N^2N_r^4}\sum_{\mathbf{k}\in\mathcal{K}^{2,4}}\boldsymbol{B}_{\mathbf{k}}^{2,4} \\
&= \frac{1}{N^2N_r^4}\left[N^2N_r^4\,\mathbb{E}\left[\chi\right]^2\mathbb{E}\left[\mathbf{x}_i\right]^2\mathbb{E}\left[\mathbf{x}_j\right]^2+\mathcal{O}(N^2N_r^3+NN_r^4)\right] \\
&= \mathbb{E}\left[\chi\right]^2\mathbb{E}\left[\mathbf{x}_i\right]^2\mathbb{E}\left[\mathbf{x}_j\right]^2+\mathcal{O}(N_r^{-1}) \\
&= \omega^2\boldsymbol{\mu}_i^2\boldsymbol{\mu}_j^2+\mathcal{O}(N_r^{-1}).
\end{aligned} \tag{A.57}$$

Thus,

$$\begin{aligned}
\mathrm{Var}\left[\hat{\omega}\hat{\boldsymbol{\mu}}_i\hat{\boldsymbol{\mu}}_j\right] &= \mathbb{E}\left[(\hat{\omega}\hat{\boldsymbol{\mu}}_i\hat{\boldsymbol{\mu}}_j)^2\right]-\mathbb{E}\left[\hat{\omega}\hat{\boldsymbol{\mu}}_i\hat{\boldsymbol{\mu}}_j\right]^2 \\
&= \left(\omega^2\boldsymbol{\mu}_i^2\boldsymbol{\mu}_j^2+\mathcal{O}(N_r^{-1})\right)-\left(\omega\boldsymbol{\mu}_i\boldsymbol{\mu}_j+\mathcal{O}(N_r^{-1})\right)^2 \\
&= \mathcal{O}(N_r^{-1}).
\end{aligned} \tag{A.58}$$

Equations (A.53) and (A.58) hold for each $r=1,\ldots,R$. Since $\hat{\boldsymbol{D}}_{\mathrm{SIR}}$ is a summation over $r=1,\ldots,R$,

$$\mathbb{E}\left[(\hat{\boldsymbol{D}}_{\mathrm{SIR}})_{ij}\right] = (\boldsymbol{D}_{\mathrm{SIR}})_{ij}+\mathcal{O}\left(N_{r_{\min}}^{-1}\right) \quad \text{and} \quad \mathrm{Var}\left[(\hat{\boldsymbol{D}}_{\mathrm{SIR}})_{ij}\right] = \mathcal{O}\left(N_{r_{\min}}^{-1}\right) \tag{A.59}$$

where $N_{r_{\min}}$ is the same from (39),

$$N_{r_{\min}} = \min_{1\le r\le R} N_r. \tag{A.60}$$

From (A.59), the mean squared error for each element of $\hat{\boldsymbol{D}}_{\mathrm{SIR}}$ is

$$\begin{aligned}
\mathrm{MSE}\left[(\hat{\boldsymbol{D}}_{\mathrm{SIR}})_{ij}\right] &= \mathrm{Bias}\left[(\hat{\boldsymbol{D}}_{\mathrm{SIR}})_{ij}\right]^2+\mathrm{Var}\left[(\hat{\boldsymbol{D}}_{\mathrm{SIR}})_{ij}\right] \\
&= \left(\mathbb{E}\left[(\hat{\boldsymbol{D}}_{\mathrm{SIR}})_{ij}\right]-(\boldsymbol{D}_{\mathrm{SIR}})_{ij}\right)^2+\mathrm{Var}\left[(\hat{\boldsymbol{D}}_{\mathrm{SIR}})_{ij}\right] \\
&= \left((\boldsymbol{D}_{\mathrm{SIR}})_{ij}+\mathcal{O}\left(N_{r_{\min}}^{-1}\right)-(\boldsymbol{D}_{\mathrm{SIR}})_{ij}\right)^2+\mathcal{O}\left(N_{r_{\min}}^{-1}\right) \\
&= \mathcal{O}\left(N_{r_{\min}}^{-1}\right).
\end{aligned} \tag{A.61}$$

Next, we examine how the element-wise mean squared error in (A.61) translates to errors in the eigenvalue estimates. By Corollary 8.1.6 in [30],

$$\left|\lambda_k(\boldsymbol{D}_{\mathrm{SIR}})-\lambda_k(\hat{\boldsymbol{D}}_{\mathrm{SIR}})\right| \le ||\boldsymbol{E}||_2 \tag{A.62}$$

where $\boldsymbol{E} = \boldsymbol{D}_{\text{SIR}} - \hat{\boldsymbol{D}}_{\text{SIR}}$. Since $||\cdot||_2 \leq ||\cdot||_F$,

$$\left|\lambda_k(\boldsymbol{D}_{\text{SIR}}) - \lambda_k(\hat{\boldsymbol{D}}_{\text{SIR}})\right| \;\leq\; ||\boldsymbol{E}||_F. \tag{A.63}$$

Squaring both sides and taking the expectation,

$$\mathbb{E}\left[\left(\lambda_k(\boldsymbol{D}_{\text{SIR}}) - \lambda_k(\hat{\boldsymbol{D}}_{\text{SIR}})\right)^2\right] \;\leq\; \mathbb{E}\left[||\boldsymbol{E}||_F^2\right]. \tag{A.64}$$

Consider the right-hand side

$$\begin{aligned}
\mathbb{E}\left[||\boldsymbol{E}||_F^2\right] &= \mathbb{E}\left[\sum_{i=1}^{m}\sum_{j=1}^{m}(\boldsymbol{E})_{ij}^2\right] \\
&= \sum_{i=1}^{m}\sum_{j=1}^{m}\mathbb{E}\left[(\boldsymbol{E})_{ij}^2\right] \\
&= \sum_{i=1}^{m}\sum_{j=1}^{m}\mathbb{E}\left[\left((\boldsymbol{D}_{\text{SIR}})_{ij} - (\hat{\boldsymbol{D}}_{\text{SIR}})_{ij}\right)^2\right] \\
&= \sum_{i=1}^{m}\sum_{j=1}^{m}\text{MSE}\left[(\hat{\boldsymbol{D}}_{\text{SIR}})_{ij}\right] \\
&= \mathcal{O}(N_{r_{\min}}^{-1})
\end{aligned} \tag{A.65}$$

as required. □

Appendix A.10. Proof of Theorem 9

Theorem 9. *Assume the same conditions from Theorem 8. Then, for sufficiently large N,*

$$\text{dist}\left(\text{colspan}(\boldsymbol{A}), \text{colspan}(\hat{\boldsymbol{A}})\right) \;=\; \frac{1}{\lambda_n(\boldsymbol{D}_{\text{SIR}}) - \lambda_{n+1}(\boldsymbol{D}_{\text{SIR}})}\,\mathcal{O}(N_{r_{\min}}^{-1}) \tag{42}$$

with high probability.

Proof. Recall from section 3.1.1 that $\boldsymbol{A}, \hat{\boldsymbol{A}} \in \mathbb{R}^{m\times n}$ contain the first n eigenvectors of $\boldsymbol{D}_{\text{SIR}}$ and $\hat{\boldsymbol{D}}_{\text{SIR}}$, respectively. Let $\boldsymbol{B}, \hat{\boldsymbol{B}} \in \mathbb{R}^{m\times(m-n)}$ contain the last $m-n$ eigenvectors of each matrix.

By Corollary 8.1.11 in [30], if

$$||\boldsymbol{E}||_2 \;\leq\; \frac{\lambda_n(\boldsymbol{D}_{\text{SIR}}) - \lambda_{n+1}(\boldsymbol{D}_{\text{SIR}})}{5}, \tag{A.66}$$

then

$$\text{dist}\left(\text{colspan}(\boldsymbol{A}), \text{colspan}(\hat{\boldsymbol{A}})\right) \;\leq\; \frac{4}{\lambda_n(\boldsymbol{D}_{\text{SIR}}) - \lambda_{n+1}(\boldsymbol{D}_{\text{SIR}})}||\boldsymbol{E}_{21}||_2 \quad \text{(A.67)}$$

where $\boldsymbol{E} = \boldsymbol{D}_{\text{SIR}} - \hat{\boldsymbol{D}}_{\text{SIR}}$ and $\boldsymbol{E}_{21} = \boldsymbol{B}^T\boldsymbol{E}\boldsymbol{A}$. We wish to show that the condition for Corollary 8.1.11 holds with high probability for sufficiently large N. Theorem 2.6 from [37] states that for any $\tau > 0$

$$\mathbb{P}\left(\, ||\boldsymbol{E}||_F \geq \tau\, \mathbb{E}\left[||\boldsymbol{E}||_F^2\right] \right) \;\leq\; \frac{1}{\tau^2} \quad \text{(A.68)}$$

which is derived from the Chebyshev inequality.

Choose τ_* to be large such that $1/\tau_*^2$ is arbitrarily close to zero. Equation (A.65) states that $\mathbb{E}\left[||\boldsymbol{E}||_F^2\right] = \mathcal{O}(N_{r_{\min}}^{-1})$. Additionally, $N_{r_{\min}} \to \infty$ as $N \to \infty$ since there is a nonzero probability of $y \in J_r$ for each $r = 1, \ldots, R$. Therefore, there exists N_* such that when $N > N_*$

$$\tau_*\, \mathbb{E}\left[||\boldsymbol{E}||_F^2\right] \;\leq\; \frac{\lambda_n(\boldsymbol{D}_{\text{SIR}}) - \lambda_{n+1}(\boldsymbol{D}_{\text{SIR}})}{5} \quad \text{(A.69)}$$

which implies

$$\mathbb{P}\left(\, ||\boldsymbol{E}||_F \leq \frac{\lambda_n(\boldsymbol{D}_{\text{SIR}}) - \lambda_{n+1}(\boldsymbol{D}_{\text{SIR}})}{5} \right) \;\geq\; 1 - \frac{1}{\tau_*^2}. \quad \text{(A.70)}$$

Since $||\cdot||_2 \leq ||\cdot||_F$,

$$\mathbb{P}\left(\, ||\boldsymbol{E}||_2 \leq \frac{\lambda_n(\boldsymbol{D}_{\text{SIR}}) - \lambda_{n+1}(\boldsymbol{D}_{\text{SIR}})}{5} \right) \geq 1 - \frac{1}{\tau_*^2}. \quad \text{(A.71)}$$

By Corollary 8.1.11 in [30],

$$\begin{aligned}
\text{dist}\left(\text{colspan}(\boldsymbol{A}), \text{colspan}(\hat{\boldsymbol{A}})\right) &\leq \frac{4}{\lambda_n(\boldsymbol{D}_{\text{SIR}}) - \lambda_{n+1}(\boldsymbol{D}_{\text{SIR}})}||\boldsymbol{E}_{21}||_2 \\
&\leq \frac{4}{\lambda_n(\boldsymbol{D}_{\text{SIR}}) - \lambda_{n+1}(\boldsymbol{D}_{\text{SIR}})}||\boldsymbol{E}||_F \\
&\leq \frac{4}{\lambda_n(\boldsymbol{D}_{\text{SIR}}) - \lambda_{n+1}(\boldsymbol{D}_{\text{SIR}})}\tau_*\mathbb{E}\left[||\boldsymbol{E}||_F^2\right] \\
&= \frac{1}{\lambda_n(\boldsymbol{D}_{\text{SIR}}) - \lambda_{n+1}(\boldsymbol{D}_{\text{SIR}})}\mathcal{O}(N_{r_{\min}}^{-1})
\end{aligned} \quad \text{(A.72)}$$

with probability $1 - 1/\tau_*^2$. □

Appendix A.11. Proof of Lemma 3.2

A vector $\mathbf{w}$ is in the null space of $\boldsymbol{C}_{\text{SAVE}}$ if and only if $\boldsymbol{\Sigma}(y)\mathbf{w} = \mathbf{w}$ for all y in the range of f.

Proof. For any $\mathbf{w} \in \mathbb{R}^m$,

$$\mathbf{w}^T \boldsymbol{C}_{\text{SAVE}} \mathbf{w} \;=\; \mathbf{w}^T \left(\int \left[\boldsymbol{I} - \boldsymbol{\Sigma}(y)\right]^2 \, d\gamma(y) \right) \mathbf{w} \;=\; \int \left(\left[\boldsymbol{I} - \boldsymbol{\Sigma}(y)\right] \mathbf{w} \right)^T \left(\left[\boldsymbol{I} - \boldsymbol{\Sigma}(y)\right] \mathbf{w} \right) \, d\gamma(y). \tag{A.73}$$

Assume $\mathbf{w} \in \text{null}(\boldsymbol{C}_{\text{SAVE}})$. Then $\mathbf{w}^T \boldsymbol{C}_{\text{SAVE}} \mathbf{w} = 0$ which means

$$\int \left(\left[\boldsymbol{I} - \boldsymbol{\Sigma}(y)\right] \mathbf{w} \right)^T \left(\left[\boldsymbol{I} - \boldsymbol{\Sigma}(y)\right] \mathbf{w} \right) \, d\gamma(y) \;=\; 0. \tag{A.74}$$

The integrand is the summation of the squares of the elements of $\left[\boldsymbol{I} - \boldsymbol{\Sigma}(y)\right] \mathbf{w}$. Thus, $\left[\boldsymbol{I} - \boldsymbol{\Sigma}(y)\right] \mathbf{w} = \mathbf{0}$, or equivalently $\boldsymbol{\Sigma}(y)\mathbf{w} = \mathbf{w}$, for all y in the range of f.

Next, assume $\boldsymbol{\Sigma}(y)\mathbf{w} = \mathbf{w}$ for all y in the range of f. Then,

$$\int \left(\left[\boldsymbol{I} - \boldsymbol{\Sigma}(y)\right] \mathbf{w} \right)^T \left(\left[\boldsymbol{I} - \boldsymbol{\Sigma}(y)\right] \mathbf{w} \right) \, d\gamma(y) \;=\; 0. \tag{A.75}$$

This implies that $\mathbf{w}^T \boldsymbol{C}_{\text{SAVE}} \mathbf{w} = 0$ which means $\mathbf{w} \in \text{null}(\boldsymbol{C}_{\text{SAVE}})$. □

Appendix A.12. Proof of Theorem 10

Theorem 10. *Given an elliptically symmetric input probability density function $\rho : \mathbb{R}^m \to \mathbb{R}^+$, assume $f : \mathbb{R}^m \to \mathbb{R}$ admits a central ridge subspace $\mathcal{S}_{f(\mathbf{x})}$. Then,* $\text{colspan}(\boldsymbol{C}_{\text{SAVE}}) \subseteq \mathcal{S}_{f(\mathbf{x})}$.

Proof. Consider the orthogonal $m \times m$ matrix $\begin{bmatrix} \boldsymbol{A} & \boldsymbol{B} \end{bmatrix}$ where $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ contains a basis for $\mathcal{S}_{f(\mathbf{x})}$ and $\boldsymbol{B} \in \mathbb{R}^{m \times (m-n)}$ is a basis for the orthogonal complement of $\mathcal{S}_{f(\mathbf{x})}$. We can decompose $\mathbf{x}$,

$$\mathbf{x} \;=\; \boldsymbol{A}\mathbf{u} + \boldsymbol{B}\mathbf{v} \tag{A.76}$$

where $\mathbf{u} = \boldsymbol{A}^T \mathbf{x} \in \mathbb{R}^n$ and $\mathbf{v} = \boldsymbol{B}^T \mathbf{x} \in \mathbb{R}^{m-n}$.

From (44),

$$\boldsymbol{\Sigma}(y) \;=\; \int \mathbf{x}\mathbf{x}^T \, d\mu_{\mathbf{x}|y}(\mathbf{x}) - \boldsymbol{\mu}(y)\boldsymbol{\mu}(y)^T. \tag{A.77}$$

The first term of (A.77) may be decomposed into

$$\int \mathbf{x}\mathbf{x}^T \, d\mu_{\mathbf{x}|y}(\mathbf{x}) \;=\; \int \left[\int \mathbf{x}\mathbf{x}^T \, d\mu_{\mathbf{x}|\mathbf{u}}(\mathbf{x}) \right] d\mu_{\mathbf{u}|y}(\mathbf{u}) \tag{A.78}$$

where $\mu_{\mathbf{x}|\mathbf{u}}$ and $\mu_{\mathbf{u}|y}$ are the same conditional measures from (33). We can rewrite the inner integral of (A.78) as

$$\int \mathbf{x}\mathbf{x}^T \, d\mu_{\mathbf{x}|\mathbf{u}}(\mathbf{x}) \;=\; \int (\boldsymbol{A}\mathbf{u} + \boldsymbol{B}\mathbf{v}) (\boldsymbol{A}\mathbf{u} + \boldsymbol{B}\mathbf{v})^T \, d\nu_{\mathbf{v}|\mathbf{u}}(\mathbf{v}) \tag{A.79}$$

$\nu_{\mathbf{v}|\mathbf{u}}$ is the conditional measure on $\mathbf{v} \in \mathbb{R}^{m-n}$ induced by $\mathbf{v} = \boldsymbol{B}^T\mathbf{x}$ and $\rho(\mathbf{x})$ over the set $\mathcal{X}(\mathbf{u})$ from (A.41). Expanding and simplifying (A.79),

$$\begin{aligned}
\int \mathbf{x}\mathbf{x}^T \, d\mu_{\mathbf{x}|\mathbf{u}}(\mathbf{x}) &= \int (\boldsymbol{A}\mathbf{u} + \boldsymbol{B}\mathbf{v}) (\boldsymbol{A}\mathbf{u} + \boldsymbol{B}\mathbf{v})^T \, d\nu_{\mathbf{v}|\mathbf{u}}(\mathbf{v}) \\
&= \int \left(\boldsymbol{A}\mathbf{u}\mathbf{u}^T\boldsymbol{A}^T + \boldsymbol{A}\mathbf{u}\mathbf{v}^T\boldsymbol{B}^T + \boldsymbol{B}\mathbf{v}\mathbf{u}^T\boldsymbol{A}^T + \boldsymbol{B}\mathbf{v}\mathbf{v}^T\boldsymbol{B}^T \right) d\nu_{\mathbf{v}|\mathbf{u}}(\mathbf{v}) \\
&= \boldsymbol{A}\mathbf{u}\mathbf{u}^T\boldsymbol{A}^T \left(\int d\nu_{\mathbf{v}|\mathbf{u}}(\mathbf{v}) \right) + \boldsymbol{A}\mathbf{u} \left(\int \mathbf{v}^T \, d\nu_{\mathbf{v}|\mathbf{u}}(\mathbf{v}) \right) \boldsymbol{B}^T \ldots \\
&\quad + \boldsymbol{B} \left(\int \mathbf{v} \, d\nu_{\mathbf{v}|\mathbf{u}}(\mathbf{v}) \right) \mathbf{u}^T\boldsymbol{A}^T + \boldsymbol{B} \left(\int \mathbf{v}\mathbf{v}^T \, d\nu_{\mathbf{v}|\mathbf{u}}(\mathbf{v}) \right) \boldsymbol{B}^T \\
&= \boldsymbol{A}\mathbf{u}\mathbf{u}^T\boldsymbol{A}^T + \boldsymbol{B}\boldsymbol{B}^T.
\end{aligned} \tag{A.80}$$

Plugging this result into (A.78),

$$\begin{aligned}
\int \mathbf{x}\mathbf{x}^T \, d\mu_{\mathbf{x}|y}(\mathbf{x}) &= \int \left[\int \mathbf{x}\mathbf{x}^T \, d\mu_{\mathbf{x}|\mathbf{u}}(\mathbf{x}) \right] d\mu_{\mathbf{u}|y}(\mathbf{u}) \\
&= \int \left[\boldsymbol{A}\mathbf{u}\mathbf{u}^T\boldsymbol{A}^T + \boldsymbol{B}\boldsymbol{B}^T \right] d\mu_{\mathbf{u}|y}(\mathbf{u}) \\
&= \boldsymbol{A} \left(\int \mathbf{u}\mathbf{u}^T \, d\mu_{\mathbf{u}|y}(\mathbf{u}) \right) \boldsymbol{A}^T + \boldsymbol{B}\boldsymbol{B}^T.
\end{aligned} \tag{A.81}$$

Let $\mathbf{w} \in \text{null}(\boldsymbol{A}^T)$. Then,

$$\begin{aligned}
\boldsymbol{\Sigma}(y)\mathbf{w} &= \left(\int \mathbf{x}\mathbf{x}^T \, d\mu_{\mathbf{x}|y}(\mathbf{x}) \right) \mathbf{w} - \boldsymbol{\mu}(y)\boldsymbol{\mu}(y)^T\mathbf{w} \\
&= \left(\int \mathbf{x}\mathbf{x}^T \, d\mu_{\mathbf{x}|y}(\mathbf{x}) \right) \mathbf{w}
\end{aligned} \tag{A.82}$$

since $\boldsymbol{\mu}(y)^T\mathbf{w} = 0$ as shown in the proof of Theorem 7 in section Appendix A.8.

Therefore,

$$\begin{aligned}
\boldsymbol{\Sigma}(y)\mathbf{w} &= \left(\int \mathbf{x}\mathbf{x}^T \, d\mu_{\mathbf{x}|y}(\mathbf{x})\right)\mathbf{w} \\
&= \left(\boldsymbol{A}\left(\int \mathbf{u}\mathbf{u}^T \, d\mu_{\mathbf{u}|y}(\mathbf{u})\right)\boldsymbol{A}^T + \boldsymbol{B}\boldsymbol{B}^T\right)\mathbf{w} \\
&= \boldsymbol{A}\left(\int \mathbf{u}\mathbf{u}^T \, d\mu_{\mathbf{u}|y}(\mathbf{u})\right)\left(\boldsymbol{A}^T\mathbf{w}\right) + \boldsymbol{B}\boldsymbol{B}^T\mathbf{w} \\
&= \mathbf{0} + \mathbf{w} \\
&= \mathbf{w}.
\end{aligned} \tag{A.83}$$

Thus, $\boldsymbol{\Sigma}(y)\mathbf{w} = \mathbf{w}$. By Lemma 3.2, $\mathbf{w} \in \text{null}(\boldsymbol{C}_{\text{SAVE}})$ and $\text{null}(\boldsymbol{A}^T) \subseteq \text{null}(\boldsymbol{C}_{\text{SAVE}})$. Since $\mathcal{S}_{f(\mathbf{x})} = \text{colspan}(\boldsymbol{A})$ is the orthogonal complement of $\text{null}(\boldsymbol{A}^T)$,

$$\text{colspan}(\boldsymbol{C}_{\text{SAVE}}) \subseteq \mathcal{S}_{f(\mathbf{x})} \tag{A.84}$$

as required. □

Appendix A.13. Proof of Theorem 11

Theorem 11. *Assume that Algorithm 2 has been applied to the data set* $\{[\, y_i \,,\, \mathbf{x}_i^T \,]\}$, *with* $i = 1, \ldots, N$, *where the* $\mathbf{x}_i$ *are drawn independently according to* $\rho(\mathbf{x})$ *and* $y_i = f(\mathbf{x}_i)$ *are point evaluations of* $f(\mathbf{x})$. *Then, for* $k = 1, \ldots, m$,

$$\mathbb{E}\left[\left(\lambda_k(\boldsymbol{D}_{\text{SAVE}}) - \lambda_k(\hat{\boldsymbol{D}}_{\text{SAVE}})\right)^2\right] \;=\; \mathcal{O}(N_{r_{\min}}^{-1}) \tag{53}$$

where $\lambda_k(\cdot)$ *denotes the* k*th eigenvalue of the given matrix.*

Proof. Recall the SAVE matrix from (50)

$$\boldsymbol{D}_{\text{SAVE}} \;=\; \sum_{r=1}^{R} \omega(r)\left[\boldsymbol{I} - \boldsymbol{\Sigma}(r)\right]^2 \tag{A.85}$$

where $\omega(r)$ is the probability mass function from (36) and $\boldsymbol{\Sigma}(r)$ is a measure of the spread of $\mathbf{x}$ over the inverse image $f^{-1}(h^{-1}(r))$ from (51). Algorithm 2 approximates $\boldsymbol{D}_{\text{SAVE}}$ using

$$\hat{\boldsymbol{D}}_{\text{SAVE}} \;=\; \sum_{r=1}^{R} \hat{\omega}(r)\left[\boldsymbol{I} - \hat{\boldsymbol{\Sigma}}(r)\right]^2 \tag{A.86}$$

where $\hat{\omega}(r)$ and $\hat{\boldsymbol{\Sigma}}(r)$ are sample estimates of $\omega(r)$ and $\boldsymbol{\Sigma}(r)$, respectively. Define $\mathcal{I}_r$ to be the set of indices such that $y_i \in J_r$, and let N_r be the cardinality of $\mathcal{I}_r$. In the SAVE algorithm, these estimates are

$$\hat{\omega}(r) \;=\; \frac{N_r}{N} \qquad \text{and} \qquad \hat{\boldsymbol{\Sigma}}(r) \;=\; \frac{1}{N_r - 1} \sum_{i \in \mathcal{I}_r} \left(\mathbf{x}_i - \hat{\boldsymbol{\mu}}(r)\right) \left(\mathbf{x}_i - \hat{\boldsymbol{\mu}}(r)\right)^T \tag{A.87}$$

where $\hat{\boldsymbol{\mu}}(r)$ is the sample estimate of the average from (10). We may rewrite $\hat{\omega}(r)$ and $\hat{\boldsymbol{\Sigma}}(r)$ as

$$\hat{\omega}(r) \;=\; \frac{1}{N} \sum_{i=1}^{N} \chi\left(y_i \in J_r\right) \quad \text{and} \quad \hat{\boldsymbol{\Sigma}}(r) \;=\; \frac{1}{N_r - 1} \sum_{i \in \mathcal{I}_r} \mathbf{x}_i \mathbf{x}_i^T - \frac{N_r}{N_r - 1} \hat{\boldsymbol{\mu}}(r) \hat{\boldsymbol{\mu}}(r)^T \tag{A.88}$$

where $\chi\left(y_i \in J_r\right)$ is an indicator function which is 1 when $y_i \in J_r$ and 0 otherwise.

We want to examine the mean squared error of the elements of $\hat{\boldsymbol{D}}_{\text{SAVE}}$. To do so, we first consider a single term in the summation from (A.86). Fix r and let

$$\begin{aligned} \omega \;&=\; \omega(r), \quad \hat{\omega} \;=\; \hat{\omega}(r), \quad \chi_i \;=\; \chi\left(y_i \in J_r\right), \\ \boldsymbol{\mu} \;&=\; \boldsymbol{\mu}(r), \quad \hat{\boldsymbol{\mu}} \;=\; \hat{\boldsymbol{\mu}}(r), \quad \boldsymbol{\Sigma} \;=\; \boldsymbol{\Sigma}(r), \quad \text{and} \quad \hat{\boldsymbol{\Sigma}} \;=\; \hat{\boldsymbol{\Sigma}}(r). \end{aligned} \tag{A.89}$$

Assume without loss of generality that $\mathcal{I}_r = \{1, \ldots, N_r\}$. This simplifies notation as we compute the mean squared error.

To compute the mean squared error, we focus on the computation of a single element in $\hat{\boldsymbol{D}}_{\text{SAVE}}$. To do this, we move the sample index to the superscript and let the subscript denote the vector or matrix element similar to section

Appendix A.9. Thus, for $1 \leq i, j \leq m$,

$$\begin{aligned}
\mathbb{E}\left[\hat{\omega}\left(\delta_{ij} - (\hat{\mathbf{\Sigma}})_{ij}\right)^2\right] &= \mathbb{E}\left[\left(\frac{1}{N}\sum_{k_1=1}^{N}\chi^{k_1}\right)\left(\delta_{ij} - \left(\frac{1}{N_r - 1}\sum_{k_2=1}^{N_r}\mathbf{x}_i^{k_2}\mathbf{x}_j^{k_2} \ \dots \right.\right.\right. \\
&\qquad \left.\left.\left. - \left(\frac{1}{N_r}\sum_{k_3=1}^{N_r}\mathbf{x}_i^{k_3}\right)\left(\frac{1}{N_r}\sum_{k_4=1}^{N_r}\mathbf{x}_j^{k_4}\right)\right)\right)^2\right] \\
&= \frac{\delta_{ij}}{N}\sum_{k_1=1}^{N}\mathbb{E}\left[\chi^{k_1}\right] - \frac{2\delta_{ij}}{N(N_r-1)}\sum_{k_1=1}^{N}\sum_{k_2=1}^{N_r}\mathbb{E}\left[\chi^{k_1}\mathbf{x}_i^{k_2}\mathbf{x}_j^{k_2}\right]\dots \\
&\quad + \frac{2\delta_{ij}}{NN_r^2}\sum_{k_1=1}^{N}\sum_{k_2=1}^{N_r}\sum_{k_3=1}^{N_r}\mathbb{E}\left[\chi^{k_1}\mathbf{x}_i^{k_2}\mathbf{x}_j^{k_3}\right]\dots \\
&\quad + \frac{1}{N(N_r-1)^2}\sum_{k_1=1}^{N}\sum_{k_2=1}^{N_r}\sum_{k_3=1}^{N_r}\mathbb{E}\left[\chi^{k_1}\mathbf{x}_i^{k_2}\mathbf{x}_j^{k_2}\mathbf{x}_i^{k_3}\mathbf{x}_j^{k_3}\right]\dots \\
&\quad - \frac{2}{NN_r^2(N_r-1)}\sum_{k_1=1}^{N}\sum_{k_2=1}^{N_r}\sum_{k_3=1}^{N_r}\sum_{k_4=1}^{N_r}\mathbb{E}\left[\chi^{k_1}\mathbf{x}_i^{k_2}\mathbf{x}_j^{k_2}\mathbf{x}_i^{k_3}\mathbf{x}_j^{k_4}\right]\dots \\
&\quad + \frac{1}{NN_r^4}\sum_{k_1=1}^{N}\sum_{k_2=1}^{N_r}\sum_{k_3=1}^{N_r}\sum_{k_4=1}^{N_r}\sum_{k_5=1}^{N_r}\mathbb{E}\left[\chi^{k_1}\mathbf{x}_i^{k_2}\mathbf{x}_j^{k_3}\mathbf{x}_i^{k_4}\mathbf{x}_j^{k_5}\right].
\end{aligned} \tag{A.90}$$

Equations (A.1) and (A.2) allow us to rewrite (A.90) as sums over various tensors. Since we fixed r, we drop the argument from the notation of (A.1) and (A.2). Thus,

$$\begin{aligned}
\mathbb{E}\left[\hat{\omega}\left(\delta_{ij} - (\hat{\mathbf{\Sigma}})_{ij}\right)^2\right] &= \frac{\delta_{ij}}{N}\sum_{k_1=1}^{N}\mathbb{E}\left[\chi^{k_1}\right] - \frac{2\delta_{ij}}{N(N_r-1)}\sum_{\mathbf{k}\in\mathcal{K}^{1,1}}\boldsymbol{B}_{\mathbf{k}}^{1,1}\dots \\
&\quad + \frac{2\delta_{ij}}{NN_r^2}\sum_{\mathbf{k}\in\mathcal{K}^{1,2}}\boldsymbol{B}_{\mathbf{k}}^{1,2} + \frac{1}{N(N_r-1)^2}\sum_{\mathbf{k}\in\mathcal{K}^{1,2}}\boldsymbol{B}_{\mathbf{k}}^{1,2}\dots \\
&\quad - \frac{2}{NN_r^2(N_r-1)}\sum_{\mathbf{k}\in\mathcal{K}^{1,3}}\boldsymbol{B}_{\mathbf{k}}^{1,3} + \frac{1}{NN_r^4}\sum_{\mathbf{k}\in\mathcal{K}^{1,4}}\boldsymbol{B}_{\mathbf{k}}^{1,4}.
\end{aligned} \tag{A.91}$$

By Lemma Appendix A,

$$\begin{aligned}
\mathbb{E}\left[\hat{\omega}\left(\delta_{ij}-(\hat{\boldsymbol{\Sigma}})_{ij}\right)^2\right] &= \frac{\delta_{ij}}{N}\sum_{k_1=1}^{N}\mathbb{E}\left[\chi^{k_1}\right]-\frac{2\delta_{ij}}{N(N_r-1)}\sum_{\mathbf{k}\in\mathcal{K}^{1,1}}\boldsymbol{B}_{\mathbf{k}}^{1,1}\dots \\
&+\frac{2\delta_{ij}}{NN_r^2}\sum_{\mathbf{k}\in\mathcal{K}^{1,2}}\boldsymbol{B}_{\mathbf{k}}^{1,2}+\frac{1}{N(N_r-1)^2}\sum_{\mathbf{k}\in\mathcal{K}^{1,2}}\boldsymbol{B}_{\mathbf{k}}^{1,2}\dots \\
&-\frac{2}{NN_r^2(N_r-1)}\sum_{\mathbf{k}\in\mathcal{K}^{1,3}}\boldsymbol{B}_{\mathbf{k}}^{1,3}+\frac{1}{NN_r^4}\sum_{\mathbf{k}\in\mathcal{K}^{1,4}}\boldsymbol{B}_{\mathbf{k}}^{1,4} \\
&= \frac{\delta_{ij}}{N}\left[N\,\mathbb{E}\left[\chi\right]\right]-\frac{2\delta_{ij}}{N(N_r-1)}\left[NN_r\,\mathbb{E}\left[\chi\right]\mathbb{E}\left[\mathbf{x}_i\mathbf{x}_j\right]+\mathcal{O}(N+N_r)\right]\dots \\
&+\frac{2\delta_{ij}}{NN_r^2}\left[NN_r^2\,\mathbb{E}\left[\chi\right]\mathbb{E}\left[\mathbf{x}_i\right]\mathbb{E}\left[\mathbf{x}_j\right]+\mathcal{O}(NN_r+N_r^2)\right]\dots \\
&+\frac{1}{N(N_r-1)^2}\left[NN_r^2\,\mathbb{E}\left[\chi\right]\mathbb{E}\left[\mathbf{x}_i\mathbf{x}_j\right]^2+\mathcal{O}(NN_r+N_r^2)\right]\dots \\
&-\frac{2}{NN_r^2(N_r-1)}\left[NN_r^3\,\mathbb{E}\left[\chi\right]\mathbb{E}\left[\mathbf{x}_i\mathbf{x}_j\right]\mathbb{E}\left[\mathbf{x}_i\right]\mathbb{E}\left[\mathbf{x}_j\right]+\mathcal{O}(NN_r^2+N_r^3)\right]\dots \\
&+\frac{1}{NN_r^4}\left[NN_r^4\,\mathbb{E}\left[\chi\right]\mathbb{E}\left[\mathbf{x}_i\right]^2\mathbb{E}\left[\mathbf{x}_j\right]^2+\mathcal{O}(NN_r^3+N_r^4)\right] \\
&= \delta_{ij}\mathbb{E}\left[\chi\right]-2\delta_{ij}\mathbb{E}\left[\chi\right]\mathbb{E}\left[\mathbf{x}_i\mathbf{x}_j\right]+2\delta_{ij}\mathbb{E}\left[\chi\right]\mathbb{E}\left[\mathbf{x}_i\right]\mathbb{E}\left[\mathbf{x}_j\right]\dots \\
&+\mathbb{E}\left[\chi\right]\mathbb{E}\left[\mathbf{x}_i\mathbf{x}_j\right]^2-2\,\mathbb{E}\left[\chi\right]\mathbb{E}\left[\mathbf{x}_i\mathbf{x}_j\right]\mathbb{E}\left[\mathbf{x}_i\right]\mathbb{E}\left[\mathbf{x}_j\right]\dots \\
&+\mathbb{E}\left[\chi\right]\mathbb{E}\left[\mathbf{x}_i\right]^2\mathbb{E}\left[\mathbf{x}_j\right]^2+\mathcal{O}(N_r^{-1}) \\
&= \mathbb{E}\left[\chi\right]\left(\delta_{ij}-\left(\mathbb{E}\left[\mathbf{x}_i\mathbf{x}_j\right]-\mathbb{E}\left[\mathbf{x}_i\right]\mathbb{E}\left[\mathbf{x}_j\right]\right)\right)^2+\mathcal{O}(N_r^{-1}) \\
&= \omega\left(\delta_{ij}-(\hat{\boldsymbol{\Sigma}})_{ij}\right)^2+\mathcal{O}(N_r^{-1}).
\end{aligned} \tag{A.92}$$

Next, we compute the variance using the property

$$\mathrm{Var}\left[\hat{\omega}\left(\delta_{ij}-(\hat{\boldsymbol{\Sigma}})_{ij}\right)^2\right]=\mathbb{E}\left[\left(\hat{\omega}\left(\delta_{ij}-(\hat{\boldsymbol{\Sigma}})_{ij}\right)^2\right)^2\right]-\mathbb{E}\left[\hat{\omega}\left(\delta_{ij}-(\hat{\boldsymbol{\Sigma}})_{ij}\right)^2\right]^2. \tag{A.93}$$

To find $\mathbb{E}\left[\left(\hat{\omega}\left(\delta_{ij}-(\hat{\boldsymbol{\Sigma}})_{ij}\right)^2\right)^2\right]$,

$$\mathbb{E}\left[\left(\hat{\omega}\left(\delta_{ij}-(\hat{\boldsymbol{\Sigma}})_{ij}\right)^2\right)^2\right] = \mathbb{E}\left[\left(\frac{1}{N}\sum_{k_1=1}^{N}\chi^{k_1}\right)^2\left(\delta_{ij}-\left(\frac{1}{N_r-1}\sum_{k_2=1}^{N_r}\mathbf{x}_i^{k_2}\mathbf{x}_j^{k_2}\ \ldots\right.\right.\right.$$
$$\left.\left.\left.-\left(\frac{1}{N_r}\sum_{k_3=1}^{N_r}\mathbf{x}_i^{k_3}\right)\left(\frac{1}{N_r}\sum_{k_4=1}^{N_r}\mathbf{x}_j^{k_4}\right)\right)\right)^4\right]. \tag{A.94}$$

By expanding, we can rewrite (A.94) in a form which can be simplified using the tensor summation notation from (A.1) and (A.2). We then apply Lemma Appendix A to simplify these summations and obtain

$$\mathbb{E}\left[\left(\hat{\omega}\left(\delta_{ij}-(\hat{\boldsymbol{\Sigma}})_{ij}\right)^2\right)^2\right] = \left(\omega\left(\delta_{ij}-(\boldsymbol{\Sigma})_{ij}\right)^2\right)^2+\mathcal{O}(N_r^{-1}). \tag{A.95}$$

Thus,

$$\begin{aligned}\mathrm{Var}\left[\hat{\omega}\left(\delta_{ij}-(\hat{\boldsymbol{\Sigma}})_{ij}\right)^2\right] &= \mathbb{E}\left[\left(\hat{\omega}\left(\delta_{ij}-(\hat{\boldsymbol{\Sigma}})_{ij}\right)^2\right)^2\right]-\mathbb{E}\left[\hat{\omega}\left(\delta_{ij}-(\hat{\boldsymbol{\Sigma}})_{ij}\right)^2\right]^2\\ &= \left(\left(\omega\left(\delta_{ij}-(\boldsymbol{\Sigma})_{ij}\right)^2\right)^2+\mathcal{O}(N_r^{-1})\right)\\ &\qquad -\left(\omega\left(\delta_{ij}-(\boldsymbol{\Sigma})_{ij}\right)^2+\mathcal{O}(N_r^{-1})\right)^2\\ &= \mathcal{O}(N_r^{-1}).\end{aligned} \tag{A.96}$$

Similar to the SIR eigenvalue convergence proof in section Appendix A.9, we extend (A.92) and (A.96) from one slice to the total summation to obtain the expectation and variance for a single element of $\hat{\boldsymbol{D}}_{\mathrm{SAVE}}$,

$$\mathbb{E}\left[(\hat{\boldsymbol{D}}_{\mathrm{SAVE}})_{ij}\right] = (\boldsymbol{D}_{\mathrm{SAVE}})_{ij}+\mathcal{O}\left(N_{r_{\min}}^{-1}\right), \qquad \mathrm{Var}\left[(\hat{\boldsymbol{D}}_{\mathrm{SAVE}})_{ij}\right] = \mathcal{O}\left(N_{r_{\min}}^{-1}\right) \tag{A.97}$$

where $N_{r_{\min}}$ from (39) denotes the minimum number of samples in any one slice. From (A.97), the mean squared error for a single element of the SAVE matrix is

$$\mathrm{MSE}\left[(\hat{\boldsymbol{D}}_{\mathrm{SAVE}})_{ij}\right] = \mathcal{O}(N_{r_{\min}}^{-1}). \tag{A.98}$$

Similar to section Appendix A.9, we use this mean squared error to obtain

$$\mathbb{E}\left[||\boldsymbol{E}||_F^2\right] \;=\; \mathcal{O}(N_{r_{\min}}^{-1}) \tag{A.99}$$

where $\boldsymbol{E} = \boldsymbol{D}_{\text{SAVE}} - \hat{\boldsymbol{D}}_{\text{SAVE}}$. Combining this result with Corollary 8.1.6 in [30] yields

$$\mathbb{E}\left[\left(\lambda_k(\boldsymbol{D}_{\text{SAVE}}) - \lambda_k(\hat{\boldsymbol{D}}_{\text{SAVE}})\right)^2\right] \;=\; \mathcal{O}\left(N_{r_{\min}}^{-1}\right). \tag{A.100}$$

□

Appendix A.14. Proof of Theorem 12

Theorem 12. *Assume the same conditions from Theorem 11. Then, for sufficiently large N,*

$$\text{dist}\left(\text{colspan}(\boldsymbol{A}), \text{colspan}(\hat{\boldsymbol{A}})\right) \;=\; \frac{1}{\lambda_n(\boldsymbol{D}_{\text{SAVE}}) - \lambda_{n+1}(\boldsymbol{D}_{\text{SAVE}})}\,\mathcal{O}(N_{r_{\min}}^{-1}) \tag{54}$$

with high probability.

Proof. In section Appendix A.13, we showed that

$$\mathbb{E}\left[||\boldsymbol{E}||_F^2\right] \;=\; \mathcal{O}(N_{r_{\min}}^{-1}) \tag{A.101}$$

where $\boldsymbol{E} = \boldsymbol{D}_{\text{SAVE}} - \hat{\boldsymbol{D}}_{\text{SAVE}}$. Given this result, the proof for Theorem 12 is identical to the proof for Theorem 9 presented in section Appendix A.10. □

Acknowledgments

We thank Luis Tenorio and Don Estep for helpful discussions and insights regarding the work presented in this paper.

References

References

[1] R. C. Smith, Uncertainty Quantification: Theory, Implementation, and Applications, SIAM, Philadelphia, 2013.
URL http://bookstore.siam.org/cs12/

[2] T. Sullivan, Introduction to Uncertainty Quantification, Springer, New York, 2015.
URL https://doi.org/10.1007/978-3-319-23395-6

[3] R. Ghanem, D. Higdon, H. Owhadi, Handbook of Uncertainty Quantification, Springer International Publishing, 2016.
URL http://dx.doi.org/10.1007/978-3-319-11259-6

[4] H. Wozniakowski, Tractability and strong tractability of linear multivariate problems, Journal of Complexity 10 (1) (1994) 96–128.
URL http://dx.doi.org/10.1006/jcom.1994.1004

[5] A. Saltelli, M. Ratto, T. Andres, J. C. Francesca Campolongo, D. Galtelli, M. Saisana, S. Tarantola, Global Sensitivity Analysis. The Primer, John Wiley & Sons Ltd, England, 2008.
URL https://doi.org/10.1002/9780470725184

[6] P. G. Constantine, Active Subspaces: Emerging Ideas for Dimension Reduction in Parameter Studies, SIAM, Philadelphia, 2015.
URL http://dx.doi.org/10.1137/1.9781611973860

[7] P. G. Constantine, E. Dow, Q. Wang, Active subspace methods in theory and practice: Applications to kriging surfaces, SIAM Journal on Scientific Computing 36 (4) (2014) A1500–A1524.
URL http://dx.doi.org/10.1137/130916138

[8] R. D. Cook, Regression Graphics: Ideas for Studying Regression through Graphics, John Wiley & Sons, Inc, New York, 1998.
URL http://dx.doi.org/10.1002/9780470316931

[9] K.-C. Li, N. Duan, Regression analysis under link violation, The Annals of Statistics 17 (3) (1989) 1009–1052.
URL http://www.jstor.org/stable/2241708

[10] K. C. Li, Sliced inverse regression for dimension reduction, Journal of the American Statistical Association 86 (414) (1991) 316–327.
URL http://doi.org/10.1080/01621459.1991.10475035

[11] R. D. Cook, S. Weisberg, Sliced inverse regression for dimension reduction: comment, Journal of the American Statistical Association 86 (414) (1991) 328–332.
URL http://www.jstor.org/stable/2290566

[12] K. C. Li, On principal hessian directions for data visualization and dimension reduction: Another application of stein's lemma, Journal of the American Statistical Association 87 (420) (1992) 1025–1039.
URL http://doi.org/10.1080/01621459.1992.10476258

[13] J. Helton, F. Davis, Latin hypercube sampling and the propagation of uncertainty in analyses of complex systems, Reliability Engineering & System Safety 81 (1) (2003) 23–69.
URL https://doi.org/10.1016/S0951-8320(03)00058-9

[14] R. D. Cook, Using dimension-reduction subspaces to identify important inputs in models of physical systems, in: Proceedings of the Section on Physical and Engineering Sciences, American Statistical Association, Alexandria, VA, 1994, pp. 18–25.
URL http://users.stat.umn.edu/~rdcook/SDR/ASA94.pdf

[15] W. Li, G. Lin, B. Li, Inverse regression-based uncertainty quantification algorithms for high-dimensional models: theory and practice, Journal of Computational Physics 321 (2016) 259–278.
URL http://dx.doi.org/10.1016/j.jcp.2016.05.040

[16] A. Pinkus, Ridge Functions, Cambridge University Press, 2015.
URL https://doi.org/10.1017/CBO9781316408124

[17] M. Fornasier, K. Schnass, J. Vybiral, Learning functions of few arbitrary linear parameters in high dimensions, Foundations of Computational Mathematics 12 (2) (2012) 229–262.
URL http://dx.doi.org/10.1007/s10208-012-9115-y

[18] H. Tyagi, V. Cevher, Learning non-parametric basis independent models from point queries via low-rank methods, Applied and Computational Harmonic Analysis 37 (3) (2014) 389–412.
URL https://doi.org/10.1016/j.acha.2014.01.002

[19] P. G. Constantine, A. Eftekhari, M. B. Wakin, Computing active subspaces efficiently with gradient sketching, in: 2015 IEEE 6th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP), 2015, pp. 353–356.
URL http://dx.doi.org/10.1109/CAMSAP.2015.7383809

[20] R. D. Cook, Graphics for regressions with a binary response, Journal of the American Statistical Association 91 (435) (1996) 983–992.
URL http://doi.org/10.1080/01621459.1996.10476968

[21] R. D. Cook, On the interpretation of regression plots, Journal of the American Statistical Association 89 (425) (1994) 177–189.
URL http://doi.org/10.1080/01621459.1994.10476459

[22] X. Yin, B. Li, R. D. Cook, Successive direction extraction for estimating the central subspace in a multiple-index regression, Journal of Multivariate Analysis 99 (8) (2008) 1733–1757.
URL http://dx.doi.org/10.1016/j.jmva.2008.01.006

[23] R. D. Cook, Save: A method for dimension reduction and graphics in regression, Communications in Statistics - Theory and Methods 29 (9-10)

(2000) 2109–2121.
URL http://dx.doi.org/10.1080/03610920008832598

[24] M. L. Eaton, A characterization of spherical distributions, Journal of Multivariate Analysis 20 (2) (1986) 272–27.
URL http://dx.doi.org/10.1016/0047-259X(86)90083-7

[25] P. Hall, K.-C. Li, On almost linearity of low dimensional projections from high dimensional data, The Annals of Statistics 21 (2) (1993) 867–889.
URL http://www.jstor.org/stable/2242265

[26] R. D. Cook, L. Forzani, Likelihood-based sufficient dimension reduction, Journal of the American Statistical Association 104 (485) (2009) 197–208.
URL http://dx.doi.org/10.1198/jasa.2009.0106

[27] A. Cohen, I. Daubechies, R. DeVore, G. Kerkyacharian, D. Picard, Capturing ridge functions in high dimensions from point queries, Constructive Approximation 35 (2) (2012) 225–243.
URL https://doi.org/10.1007/s00365-011-9147-6

[28] P. G. Constantine, A. Eftekhari, R. Ward, A near-stationary subspace for ridge approximation, arXiv:1606.01929v1.
URL http://arxiv.org/abs/1606.01929

[29] R. D. Cook, C. J. Nachtsheim, Reweighting to achieve elliptically contoured covariates in regression, Journal of the American Statistical Association 89 (426) (1994) 592–599.
URL http://doi.org/10.1080/01621459.1994.10476784

[30] G. H. Golub, C. F. Van Loan, Matrix Computations, 2nd Edition, JHU Press, Baltimore, 1996.
URL https://jhupbooks.press.jhu.edu/content/matrix-computations

[31] A. Glaws, P. G. Constantine, J. Shadid, T. M. Wildey, Dimension reduction in MHD power generation models: dimensional analysis and active

subspaces, arXiv:1609.01255v1.
URL https://arxiv.org/abs/1609.01255

[32] T. Cowling, R. B. Lindsay, Magnetohydrodynamics, Physics Today 10 (1957) 40.
URL http://doi.org/10.1063/1.3060498

[33] P. G. Constantine, Z. del Rosario, G. Iaccarino, Many physical laws are ridge functions, arXiv:1605.07974.
URL https://arxiv.org/abs/1605.07974

[34] M. M. Rao, R. J. Swift, Probability Theory with Applications, Springer, New York, 2006.
URL https://doi.org/10.1007/0-387-27731-5

[35] R. L. Wheeden, A. Zygmund, Measure and Integral: An Introduction to Real Analysis, CRC, Boca Raton, 1977.
URL https://www.crcpress.com/Measure-and-Integral-An-Introduction-to-Real-Analysis-Sec
Wheeden/p/book/9781498702898

[36] M. E. Johnson, Multivariate Statistical Simulation, Wiley, Hoboken, 1987.
URL https://doi.org/10.1002/9781118150740

[37] G. W. Stewart, Stochastic perturbation theory, SIAM Review 4 (32) (1990) 579–610.
URL https://doi.org/10.1137/1032121